

Harry E. Martz, Center Leader

The Nondestructive Characterization (NDC) Center (formerly the Nondestructive Evaluation Thrust Area) at Lawrence Livermore National Laboratory (LLNL) supports initiatives that advance inspection science and technology. Our goal is to provide cutting-edge technologies that show promise for quantitative inspection, and characterization tools three to five years into the future.

The NDC Center strategic objectives involve quantitative NDC, and fast nanometer-scale imaging. This past year the NDC Center portfolio of projects focused on quantitative NDC in support of weapons system performance and laser systems.

The NDC Center supports a multidisciplinary team, consisting of mechanical and electrical engineers, physicists, material and computer scientists, chemists, technicians, and radiographers. These team members include personnel that cross departments within LLNL. Some are from academia and industry, within the United States and abroad. This collaboration brings together the necessary and diverse disciplines to provide the key scientific and technological advancements required to meet LLNL programmatic and industrial NDC challenges.

NDC provides materials characterization inspections of finished parts and complex objects, to find

flaws and fabrication defects and to determine their physical and chemical characteristics. In addition, applying NDC throughout the life cycle of a part saves time and money and improves quality. For example, NDC is being applied at the beginning of a part to develop new materials and to aid in process design and development. NDC encompasses process monitoring and control sensors and the monitoring of in-service damage. NDC is also being applied at the end of a part's life for proper reuse or safe and proper disposal. Therefore, NDC is becoming both a front-line and an end-of-line technology that strongly impacts issues of certification, life prediction, life extension, reuse, and disposal.

To meet today's programmatic demands, it is important to increase collaboration among LLNL engineering centers, departments, and programs. This year, collaborators included the Microtechnology Center, Defense Systems, Materials Science and Technology, Physics, the Life Extension Program, the National Ignition Facility Project, Computations, and Nonproliferation. Such collaborations enable us to stay at the leading edge of NDC technology, research, development, and application in support of LLNL programs.

Center for Nondestructive Characterization



5

5. Center for Nondestructive Characterization

Overview

Harry E. Martz, Center Leader

Techniques for Enhancing Laser Ultrasonic Nondestructive Evaluation

Graham H. Thomas, Robert D. Huber, Diane J. Chinn, James V. Candy, and James Spicer5-1

Nondestructive Evaluation of an Aluminum Alloy Using Hyperspectral Infrared Imaging

Randy S. Roberts5-9

In-Situ Identification of Anti-Personnel Mines Using Acoustic Resonant Spectroscopy

Randy S. Roberts and Roger L. Perry5-13

An Acoustic Technique for the Non-Invasive *In-Situ* Measurement of Crystal Size and Solution Concentration

Diane J. Chinn, Paul R. Souza, and Harry F. Robey5-19

Micro X-Ray Computed Tomography for PBX Characterization

Diane J. Chinn, Jerry J. Haskins, Clinton M. Logan, Dave L. Haupt, Scott E. Groves, John Kinney, and Amy Waters5-23

Evaluation of an Amorphous Selenium Array for Industrial X-Ray Imaging

Clinton M. Logan, Jerry J. Haskins, Kenneth E. Morales, Earl O. Updike, James M. Fugina, Anthony D. Lavietes, Daniel J. Schneberk, Gregory J. Schmid, Keo Springer, Peter Soltani, and Kenneth Swartz5-27

LANDMARC Radar Mine Detection

Stephen G. Azevedo, Jeffery E. Mast, James M. Brase, and E. Tom Rosenbury5-39

IMAN-3D: A Software Tool-Kit for 3-D Image Analysis

Sailes K. Sengupta5-51

Image Recovery Techniques for X-Ray Computed Tomography in Limited-Data Environments

Dennis M. Goodman, Jessie A. Jackson, Maurice B. Aufderheide, and Erik M. Johansson5-61



Techniques for Enhancing Laser Ultrasonic Nondestructive Evaluation

Graham H. Thomas, Robert D. Huber, and Diane J. Chinn

*Manufacturing and Materials Engineering Division
Mechanical Engineering*

James V. Candy

*Electronics Engineering Technologies Division
Electronics Engineering*

Professor James Spicer

*Johns Hopkins University
Baltimore, Maryland*

Ultrasonic nondestructive evaluation has been an extremely powerful tool for characterizing materials and detecting defects for many years. Piezoelectric transducers, traditionally used to generate and detect high-frequency acoustic energy, usually require a liquid medium to couple the ultrasound into the material being characterized. This need for a couplant restricts the applicability of ultrasonics since many materials can be damaged by the use of couplants. We are developing a technology that generates and detects the ultrasonic pulses with lasers and thus there is no requirement for couplants. The ultrasound is generated and detected in a remote, non-contact manner since only the laser light is in contact with the material. Laser-based ultrasound has wide application in many Lawrence Livermore National Laboratory (LLNL) programs, especially when remote and/or non-contact sensing is necessary.

Introduction

Although the use of ultrasound for nondestructive evaluation is a mature technology, there continue to be many advances that expand its role in material characterization, manufacturing process control, defect detection, and life cycle management. Ultrasonics is evolving with improvements such as higher frequencies to sense smaller defects, modern signal processing methods to increase sensitivity, classification algorithms for defect characterization and the most recent imaging techniques to display defects.

Despite these advances, a universal limitation of traditional piezoelectrically generated and detected ultrasound is the need to transmit the acoustic energy from the transducer into the part through a fluid, most often water. For many parts and materials, particularly those of interest to LLNL and the Department of Energy (DOE), it is extremely desirable to eliminate this couplant.

Laser generation and detection of ultrasonic energy provides a method to perform remote, non-contact ultrasonics.¹ It allows ultrasonic evaluations in high-temperature and radioactive environments, in applications where access is restricted, such as in a vacuum, and on materials that would be damaged by couplant contamination. For ultrasonic inspections on radioactive materials, any couplant used becomes hazardous waste, and thus laser ultrasonics reduces hazardous waste since no couplant is required.

Applications for laser-based ultrasound continue to be implemented as breakthroughs in the technology occur, but there is still much to be understood before its full potential can be realized. A significant limitation of laser-based ultrasound is its poor sensitivity as compared to the sensitivity of traditional piezoelectric-based ultrasonics. To improve the sensitivity, research is being pursued in the areas of improved ultrasonic generation, better detectors, and signal processing to make

laser ultrasonics viable. Also beam-forming and other signal processing techniques have been developed to improve the defect detection levels of laser acoustics. Significant improvements in both the generation and detection aspects of LLNL's laser-based ultrasound capability were realized during this project. Laser-based ultrasound will continue to have an increased role in nondestructive evaluation as the sensitivity limitation is solved through research.

Progress

This project is exploring the science of generating acoustic energy with a laser pulse, and the methodology of detecting ultrasonic signals with laser interferometers. The goal is not to replace existing piezoelectric-based ultrasonics with laser generation and detection, but to supplement the technology to expand its role in LLNL programs. The following text describes the equipment and software algorithms for laser-based ultrasonics developed by this project.

Signal and Image Processing and Beam-Forming

A large portion of our effort has been in the area of signal and image processing, an approach to improve the sensitivity of laser-based ultrasonics. It helps extract more information from the experimentally obtained data. In the past we have demonstrated the benefits of several signal processing approaches that greatly improve laser-based ultrasonic sensitivity to finding defects.

A model-based signal processing technique has been developed and tested.^{2,3} This technique predicts the acoustic signals that are generated by a specific laser. A code (WAVER) developed at the Johns Hopkins University⁴ has been modified to handle materials and laser configurations that are of interest to LLNL programs. The model-based signal processing was implemented on ultrasonic signals generated with a Nd:YAG laser and detected with a Michelson interferometer.⁵ This demonstrated technology is significant to the implementation of laser-based ultrasound for many nondestructive evaluation applications since it allows the application of laser-based ultrasonics for materials where acoustic attenuation is large.

Another signal processing approach that we have demonstrated is beam-forming. Beam-forming improves defect detection sensitivity by viewing the

defect from several directions and combining the information from each direction in a manner to accentuate the defect's image. Previous work modeled the acoustic beam from an array of senders and receivers. This year the predicted defect sensitivity was confirmed with beam-forming algorithms processing experimental data. This method enhances the detection and display of defects by combining the information from an array of sensors.^{5,6} We have modeled the array and its response to a flaw.

Based on these results, we configured a synthetic laser ultrasonic array and confirmed the model. This new technology is one of our methods for processing ultrasonic data.

This past year we have developed matched-field imaging of the laser-based ultrasound signals. This signal processing technique uses a novel correlation-canceling approach to eliminate noise, thereby increasing the signal-to-noise ratio of the experimentally obtained data. Eleven linear scans of Nd:YAG pulsed-laser-generated, laser-based Michelson-interferometer-detected data were obtained on a 9.5-mm-thick aluminum plate. Each scan consists of 21 amplitude vs time waveforms.

Figure 1 shows the source and detection locations for these scans. There were 11 source locations spaced 2 mm apart, and 21 detection locations spaced 1 mm apart. The source position was held fixed for a scan (21 different detection locations), and then moved 2 mm for the next scan. A 1/16-in. (1.6 mm)-diameter hole was then drilled into the aluminum plate to simulate a defect, and the plate was scanned over the same region that was scanned prior to the hole being drilled.




Figure 1. Source and detection sites for laser ultrasonic generation and detection on an aluminum plate with hole. Data obtained before and after the hole was drilled were used to obtain data in Figure 2.

The waveforms obtained after the hole was drilled, shown in **Fig. 2a**, were used along with the reference waveforms obtained before the hole was drilled (**Fig. 2b**). An optimal correlation-canceling scheme was developed to extract just the hole information, as shown in **Fig. 2c**, with a single channel result shown in **Fig. 2d**.

Processing yielded the canceled signal from the pre-hole and post-hole data. These signals were used to generate images, which show the effects of the presence of the hole on the ultrasound for different source and detector locations (**Fig. 3**). This demonstrates how laser-based ultrasound with signal processing can be used for defect detection and location.

Facility

Significant improvements in our laboratory have been realized over the duration of the project. At the start, LLNL had a Michelson interferometer and a small Nd:YAG laser for laser-based ultrasound work. These items had only limited application, and there was only a small area devoted to laser-based ultrasound. Once the project began, improvements to the correction circuit of the Michelson interferometer were made that greatly increased its stability. Modifications to both the electronic and mechanical components were made.

We have added a LISOR (Light In, Signal Out Receiver) interferometer to expand our detection




Figure 2. Laser ultrasonic correlation cancellation for enhanced flaw detection. Top left plot (a) shows the waveforms obtained after the hole was drilled (Measurement) for one source location. Top right signals (b) were obtained prior to drilling of the hole (Reference). Bottom left plot (c) displays the cancelled signals. Bottom right display (d) shows an individual wave for the pre- and post-drilled cases with the corresponding cancelled signal.

capabilities. This instrument is a Fabry-Perot-based system that complements the earlier path-stabilized Michelson interferometer. The Michelson requires highly reflective surfaces to sense the ultrasonic signal, whereas the Fabry-Perot works on less reflective surfaces.

This system includes an interferometer and a laser. The laser is a frequency doubled Nd:YAG, which has an output power of 200 mW at a wavelength of 532 nm. Fabry-Perot interferometers can use light scattered from rough surfaces since they can work with multiple speckles. Michelson interferometers work with a single speckle only, which limits their use to polished surfaces.

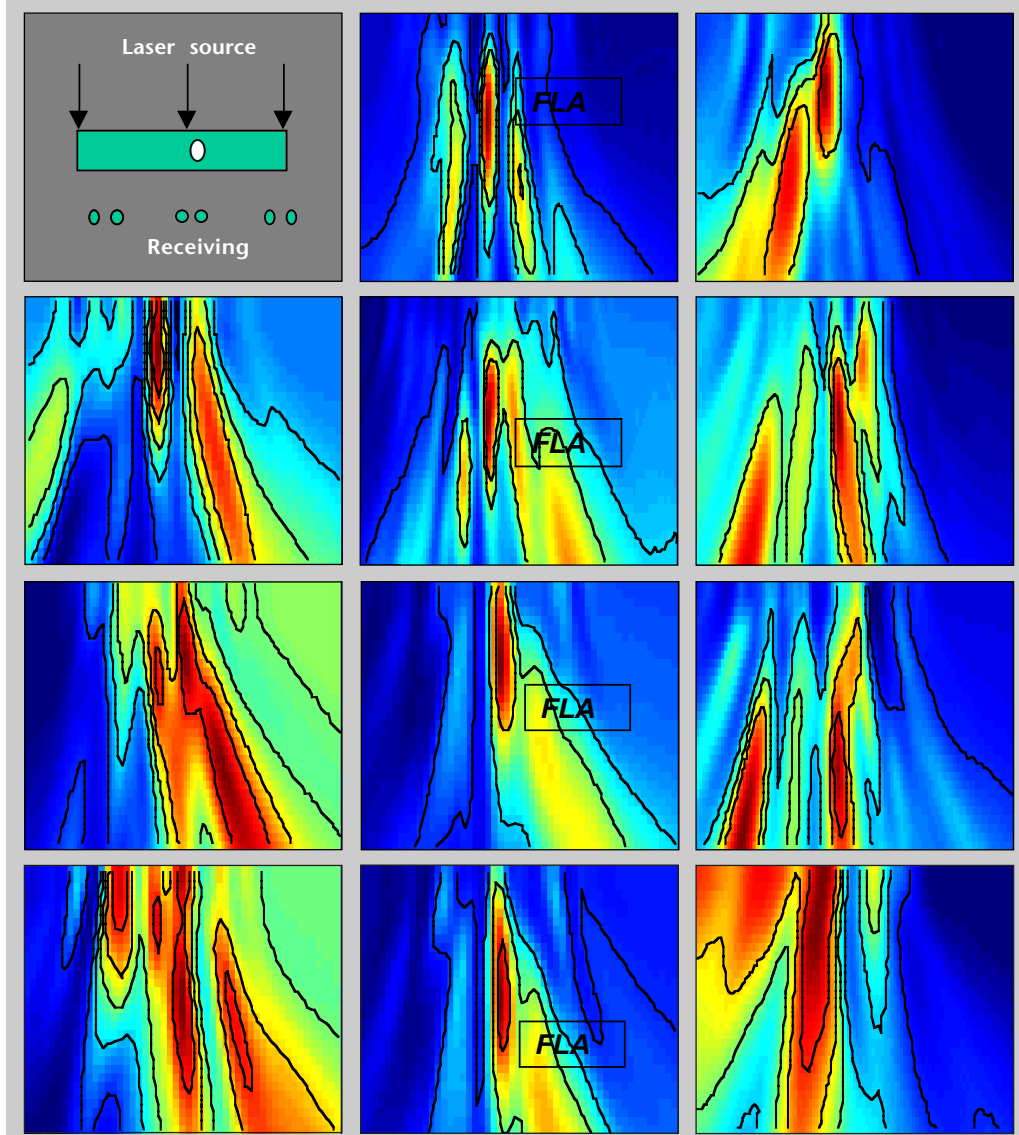
The LISOR system has been used successfully in the laboratory to detect ultrasound propagating through various materials and specimens. To

increase the flexibility of the ultrasonic detection system, an optical fiber cable 20 m long transmits the laser light from the interferometer's laser to the specimen, allowing the testing of items up to that distance away from the laser and interferometer.

Fiber optics have the advantages of allowing large distances between equipment and test location, providing access to hostile or hard to reach environments, changing test configuration easily, and eliminating the need for line of sight between laser and part. The Michelson interferometer has a frequency range of DC to 40 MHz, while the Fabry-Perot has a frequency range of 3 to 100 MHz. Together, these interferometers provide a full range of detection capabilities for most ultrasonic evaluation needs.

Another addition to the laser ultrasonic facility is a much improved source laser that provides more

Figure 3. Images obtained from correlation cancellation. The 11 images correspond to 11 different source locations for linear scans on the aluminum plate. The hole appears sharpest when the source is directly over the hole.



light energy, a higher repetition rate and a uniform spot. This Nd:YAG source laser can operate at two wavelengths and has much better beam quality than the smaller Nd:YAG laser. This laser has made a tremendous difference in the signals obtained, and is aiding greatly in the model-based processing work being performed at LLNL.

As part of the data acquisition and imaging effort, we acquired computer-controlled motion stages. These stages include two translation stages and one rotation stage that allow the movement of specimens for scanning. Data acquisition and motion control programs have been written to control the stages and capture the ultrasonic signals. The signal processing and beam-forming algorithms are combined with this software to render images of defects from the data obtained when specimens are scanned. Most nondestructive evaluation tasks require the results to be an image of the defects. Scanning capability is needed to generate ultrasonic images.

Fiber Optics

Fiber optics are a significant development for laser-based ultrasonic implementation since they allow lasers and other equipment to be located in one place, with the object under test located in another place. The laser light can then be transmitted from one location to the other, completely contained in the fibers, thus eliminating the hazards of transmitting beams through the air. Fiber optics have been developed and implemented in several applications. The first such activity involved placing the probe arm of the Michelson interferometer in a single mode optical fiber.

Figure 4 is a comparison of signals obtained using the Michelson interferometer for the two cases of 1) optical fiber probe path, and 2) no optical fiber. The two waveforms show the ultrasound detected in an aluminum plate for a through transmission case for a laser-generated thermoelastic source. The thermoelastic source causes no damage to the sample and is therefore useful for nondestructive evaluation and characterization.

The waveforms are quite similar, with only a slight decrease in the signal-to-noise for the case of the fiber, because there is some loss of light associated with the use of the optical fiber. Although there is a slight decrease in the signal-to-noise ratio for the case of optical fiber in the interferometer, there may be situations where the benefits of fiber outweigh the decrease in signal-to-noise.

Fiber optics are also used with the Fabry-Perot interferometer. Initially the illuminating laser was transmitted through air and only the light reflecting off the specimen was transferred to the Fabry-Perot interferometer in a multi-mode fiber. Then we demonstrated that the light from the part-illuminating source laser could be transmitted through a multi-mode fiber. Thus all of the travel path of the laser light can be contained in optical fiber.

Fiber optic transmission of the laser beam provides a means to place the laser away from the part and still minimize the hazard to people. Also fiber optic light transmission allows for ultrasonic evaluations inside components with access provided by a small hole.




Figure 4. Comparison of signals. The upper waveform was captured using the Michelson interferometer with no fiber, while the lower was captured using fiber in the probe arm. Only a slight decrease in signal-to-noise ratio is seen. The above are through-transmission on epicenter.

Applications

The laser-based ultrasound laboratory has the capability to evaluate a wide variety of materials, and has looked at materials including aluminum, stainless steel, uranium and paper. Thus we have demonstrated the applicability of laser-based ultrasound on a variety of important, program-critical materials. This capability will allow us to support programs at LLNL such as those for weapons, lasers, nonproliferation, and energy. The following section briefly describes programmatic applications that resulted from this research.

Process Control. As a sensor for process control, ultrasonics is extremely valuable. The speed at which ultrasonic testing can acquire measurements facilitates feedback control for machining operations. We selected a process control problem as a vehicle to direct our research. The process was a cutting application with a state-of-the-art laser that produces very short pulses. The pulses from the cutting laser generate acoustic signals that contain information about the cutting process.

Our first challenge was to understand the phenomenon and the types of ultrasonic waves that are generated. We modeled the system with a computer code that calculates the expected ultrasonic signal, based on laser input parameters and the material properties and geometry. Initial tests were run on thin stainless steel plates.⁷ The cutting laser beam was directed at the plates, and the Michelson interferometer was used to detect the ultrasound generated by the cutting beam.

The Michelson interferometer was able to detect both bulk and surface waves, depending on the test set-up. In one set-up, detection was performed on the side of the plate opposite that on which the cutting beam was incident, for the capture of bulk waves. In the other set-up, the cutting beam and the interferometer beam were both incident on the same side, for the capture of the surface waves. This configuration represented the process monitoring set-up.

Next, the Michelson interferometer was used to detect the ultrasound generated by the cutting laser in a real part.⁵ At the start of the cut, a clear signal was seen, and this signal disappeared by the time cut-through occurred. This loss of acoustic signal at cut-through is the parameter selected for sensing when the laser had cut through the part. After successfully detecting the signals generated by the cutting laser with the Michelson interferometer, we switched to a Fabry-Perot interferometer for ultrasound detection. Whereas a strip of reflective tape

was required at the detection site when using the Michelson interferometer, no tape is needed for the Fabry-Perot, and thus detection on non-specularly reflecting surfaces is possible.

The case for the process monitoring application involves non-specularly reflecting components. The Fabry-Perot interferometer is designed to operate with minimal operator intervention. Light for the interferometer is transmitted to and from the cutting chamber by optical fibers 20 m in length. To test the sensitivity, a laser-based ultrasonic through-transmission test on an aluminum plate was run using a Nd:YAG laser to generate ultrasound in the thermoelastic regime and the Fabry-Perot interferometer for detection. The waveform for this case is shown in **Fig. 5**.

The Fabry-Perot interferometer was also used to detect laser-generated ultrasound on a stainless steel surrogate part, to more nearly mimic the testing on actual parts. **Figure 6** shows a waveform signal-averaged from 10 laser pulses. The stainless steel part was a thin shell, and the waveform shows a dispersive plate wave. Over the course of the laser cutting project, various parameters of the cutting pulses such as energy and repetition rate were changed. A piezoelectric transducer was used to verify that the pulses were still generating ultrasound once the Fabry-Perot interferometer was obtained. Once the feasibility was demonstrated, the program funded this feedback control sensor as a standby option that can be implemented if needed.

Laser-based Ultrasound Sensor for Paper Manufacturing. The DOE has an initiative to reduce energy consumption in the largest industries in the




Figure 5. Fabry-Perot-interferometer-detected signal of laser-generated ultrasound. This is a through-transmission on epicenter case in aluminum for a thermoelastic source.

United States. Paper manufacturing is one such industry. We are collaborating with Lawrence Berkeley Laboratory (LBL) on developing a laser ultrasonic sensor. Real time process control afforded by this sensor will increase the quality of the product and reduce the energy consumption.

The researchers from LBL and LLNL are working to demonstrate ultrasonic characterization of paper as it is being processed. Plate waves are generated in the paper with a laser, and a Fabry-Perot interferometer detects the ultrasonic signals. Changes in acoustic velocity and attenuation signify variations in the paper processing. This project is on-going and may result in a demonstration at a factory.

Work on Aluminum at Elevated Temperatures.

One task of a cooperative research and development agreement (CRADA) between the DOE and the "Big Three" automobile manufacturers is to explore nondestructive evaluation techniques to improve the casting of aluminum. Nondestructive evaluation could play a crucial role in improving cast parts by determining the condition of the molten materials prior to casting. Impurities in the molten metal cause defective castings. These defective castings increase manufacturing costs as well as energy consumption.

For example, if it can be determined that a particular melt contains an excessive amount of oxide, then that melt would not be used to cast parts. Laser-based ultrasound offers a means of testing metallic materials at elevated temperatures since it is a non-contact technique, whereas contact techniques, such as piezoelectric transducer ultrasound, have an upper limit on temperature, above which they will not work.




Figure 6. Fabry-Perot-interferometer-detected signal of laser-generated ultrasound in a stainless steel shell. The waveform shows the dispersion of the plate waves, that is, the high frequencies travel faster than the low frequencies.

We are funded to demonstrate laser generation and detection of ultrasound in molten metal. A furnace that has a port to pass the laser beam has been installed in the laser-based ultrasonic laboratory. Impurities in the liquid aluminum will reflect the acoustic energy that was generated by a laser. The acoustic energy reflected by impurities in the molten material is detected by an interferometer. This system will remotely monitor the quality of raw material feeding the casting process and sense problems before defective parts are produced.


Future Work

Although LLNL funding of laser-based ultrasound has produced a working facility for nondestructive characterization and evaluation of materials, there is still much work to be done in this area. Refinement of the computer code that calculates the ultrasound generated in materials from laser pulses will allow model-based processing on more complex parts, such as parts containing defects. Continued advances in the signal and image processing aspect of the work will lead to increases in the signal-to-noise ratios of experimentally obtained data, and will thus allow laser-based ultrasonic testing on a wider variety of materials and parts.

Recently, a fair amount of research has been performed in the area of higher frequency laser-based ultrasound. An interferometric technique for detection has been reported,⁸ with a number of references in laser-generated high-frequency ultrasound. This includes acoustic frequencies >100 MHz and well into the GHz range. These high frequencies are required when working with materials whose characteristic dimension or defect size is <1 μm . This relatively new area of research has the potential for meeting the high spatial resolution demands emerging in LLNL program applications.

References

1. Scruby, C. B., and L. E. Drain (1990), "Laser Ultrasonics, Techniques and Applications," Adam Hilger, New York, New York.
2. Candy, J. V., G. H. Thomas, D. J. Chinn, and J. B. Spicer (1996), "Laser Ultrasonic Signal Processing: A Model-Reference Approach," *J. Acous. Soc. Am.*, **100** (1).
3. Huber, R. D., D. J. Chinn, G. H. Thomas, J. V. Candy, and J. Spicer (1997), "Model-Based Signal Processing for Laser Ultrasonic Signal Enhancement," *Review of Progress in Quantitative Nondestructive Evaluation*, **16**, pp. 757-764, Plenum Press, New York, New York.

4. Spicer, J. (1991), "Laser Ultrasonics in Finite Structures: Comprehensive Modeling with Supporting Experiment," Ph.D. Thesis, Johns Hopkins University, Baltimore, Maryland.
5. Thomas, G. H., D. J. Chinn, R. D. Huber, J. V. Candy, and J. B. Spicer (1998), "Techniques for Enhancing Laser Ultrasonic Nondestructive Evaluation," Lawrence Livermore National Laboratory, Livermore, California (UCRL-ID-129207).
6. Johnson, D., and D. Dudgeon (1993), "Array Signal Processing: Concepts and Techniques," Prentice-Hall, Princeton, New Jersey.
7. Chinn, D. J., R. D. Huber, D. D. Scott, G. H. Thomas, J. V. Candy, and J. B. Spicer (1997), "New Techniques in Laser Ultrasonic Testing," Lawrence Livermore National Laboratory, Livermore, California (UCRL-ID-125476).
8. Fiedler, C. J. (1996), "The Interferometric Detection of Ultrafast Pulses of Laser Generated Ultrasound," Wright Laboratory (WL-TR-96-4057). 

Nondestructive Evaluation of an Aluminum Alloy Using Hyperspectral Infrared Imaging

Randy S. Roberts

*Defense Sciences Engineering Division
Electronics Engineering*

Fourier transform spectroscopy has found application in many areas, including atmospheric chemistry and material characterization. This report describes an investigation into the application of the Lawrence Livermore National Laboratory (LLNL) Hyperspectral Infrared Imaging Spectrometer (HIRIS) to the nondestructive evaluation of blocks of aluminum alloy. We describe the HIRIS system and the aluminum alloys used in the investigation, and the technique used to collect the hyperspectral imagery. We discuss the processing required to transform the data into usable form and a technique to analyze the data. We also provide some preliminary results.

Introduction

The HIRIS is a second-generation Fourier transform spectrometer developed at LLNL (see Reference 1 for an excellent discussion of HIRIS's predecessor). While it has been used primarily in remote sensing applications, the focus of this project is to assess its usefulness in nondestructive evaluation applications.² The HIRIS consists of a Michelson interferometer, a cryogenically cooled

silicon arsenide (Si:AsBIB) focal plane array, and associated optics.

Figure 1 is a block diagram of a hyperspectral imaging spectrometer.³ Infrared light enters the interferometer and is split into two paths by the beam-splitter. Along one path, a fixed mirror reflects the light back through the beam-splitter and onto the focal plane array. Along the second path, a movable mirror reflects the light back through the beam-splitter which in turn reflects it onto the focal




Figure 1. Schematic diagram of the HIRIS imaging system. Light enters the interferometer from the left, and is split into two paths by the beam-splitter: towards a fixed-position mirror, and towards a movable-position mirror. The mirrors reflect the light back to the beam-splitter, which directs the light onto the focal plane array. Translation of the movable mirror results in a time-varying interference pattern on the focal plane array.

plane array. Note that the optical length for the first path is fixed, while the length of the second path can be adjusted by moving the mirror.

If the light impinging on the focal plane array is in phase, constructive interference results. If the light is out of phase, destructive interference results. If the movable mirror translates with a fixed velocity, a time-varying interference pattern is formed on the focal plane array. Since HIRIS uses a focal plane array, a time-series of interference data, called an interferogram, evolves in each pixel. Fourier transformation of all interferograms results in the desired hyperspectral image data.

The objects of our interest are several aluminum blocks provided to LLNL by General Motors Corporation (GM). The blocks are made of an aluminum alloy denoted A356.⁴ This alloy consists primarily of aluminum (~92%), silicon (~7%) and trace amounts of copper, magnesium, manganese, iron, and zinc. The alloy has excellent corrosion resistance and casting properties, good machinability and weldability, and high strength. It is most often used in the automotive and aircraft industries where high strength and low corrosion are required.

The blocks we received from GM were subjected to compression-tension cycles until they fractured. Typically, A356 reliably fails after approximately 500,000 cycles, but these samples failed after approximately 5000 cycles. It is conjectured that the premature failure was due to the addition of salts, such as potassium chloride and sodium chloride, during processing.

The objective of our investigation is to study the fracture sites in the A356 samples using hyperspectral infrared imagery, and if possible, determine the validity of this conjecture.




Figure 2. Plot of an exemplar interferogram from an A356 sample. This interferogram was formed in pixel (64, 64), and its range has been limited to reveal the most interesting structure.

Progress

For these experiments, the HIRIS was reconfigured from a remote sensor to a laboratory instrument. Reconfiguration essentially consisted of installing a collimator lens at the input point of the interferometer. The focal length of the collimator lens was 100 mm, and the focal length of the condenser lens was 226 mm, thereby yielding a system magnification of 2.26. The focal plane array in the HIRIS is a 128×128 -pixel array with a pixel size of $75 \mu\text{m}/\text{pixel}$. The resulting resolution of the system is thus on the order of $33 \mu\text{m}/\text{pixel}$.

To collect the data, the aluminum blocks were heated with a hot air gun until they reached a temperature in excess of 50°C . The blocks were then held at this temperature for the duration of the data collection. In addition to the A356 blocks, calibration data was collected from a blackbody source. The source was set to 55°C for a hot reference, and 45°C for a cold reference.

A scan in the HIRIS consists of translating the movable mirror at a constant velocity through a fixed distance. As a point of terminology, the resulting sequence of images produced by a scan is referred to as a raw data cube. The scan time and sampling rate of the interferograms determine the spectral resolution of the hyperspectral data. For these experiments, the resulting resolution of the system was set to 16 cm^{-1} .

To increase the signal-to-noise ratio of features in the spectral data, an ensemble of data cubes were collected from each object of interest. In our case, 32 data cubes were collected from each A356 sample, and 16 data cubes each were collected from the hot and cold body sources for each A356 sample.

A portion of one of the interferograms associated with pixel (64, 64) of an A356 sample is shown in **Fig. 2**. Note the structure of the interferogram between frames 220 and 580. An image associated with this part of the data cube, frame 278, reveals a ring-like interference pattern, as illustrated in **Fig. 3**.

Interferograms, while illuminating some signal structure in the data cube, are difficult to interpret. However, the spectral information contained in interferograms can be quite revealing. To begin the processing, the data cubes were first averaged to increase the signal-to-noise ratio. As a result of the averaging, three data cubes were produced for each collection: an A356 data cube, a hot calibration data cube, and a cold calibration data cube.

The spectral content of a data cube is found by Fourier transforming each interferogram in the data cube. To reduce spectral sidelobes, each interferogram is multiplied by a triangular data-tapering window prior to Fourier transformation. To ease the computational burden of subsequent analysis, the fringe portions of the spectra are discarded. Finally, the A356 data is calibrated using the hot and cold references.¹ The calibrated data cube was then analyzed as described below.

Spectral data in a data cube is analyzed using a technique called Principle Component Analysis,^{5,6} a powerful analytical tool for analyzing complex data sets. In essence, data is rotated into a coordinate system with some highly desirable properties such as orthogonality and minimum basis restriction error. The technique begins by estimating the correlation matrix of the data. The eigenvalues and eigenvectors of the correlation matrix are then computed. By themselves, the eigenvectors form an orthogonal coordinate system in which to represent the data. However, the eigenvectors associated with the largest eigenvalues correspond to the axes that most efficiently represent the data. By representing the data in the new coordinate system, that is, a coordinate system using the eigenvectors associated with the largest eigenvalues, the most important features of the data are more readily determined.

The distribution of eigenvalues can reveal a great amount of information about the structure of the data. A plot of the logarithm of the eigenvalues of one of the A356 samples is shown in **Fig. 4**. As illustrated in the figure, the magnitude of the eigenvalues decreases rapidly over the first four or five. From approximately the fifth to perhaps the twentieth eigenvalue, the decrease is more gentle, but the eigenvalues still have a large magnitude.

The modes associated with these twenty eigenvalues comprise the main structure of the data. If the data were reconstructed from these modes, the residual error would be minimal. The modes associated with the remaining eigenvalues are not important to the analysis. In particular, the modes from twenty to approximately eight-five are associated with system noise. The remaining eigenvalues belong to degenerate modes. They are artifacts of the processing, and are not useful in the analysis. Once the important features of the data set have been determined using eigenanalysis, they can be interpreted for physical meaning. Indeed, this technique has been used to identify constituent components of mineral specimens.¹

Summary

The HIRIS developed at LLNL has found many uses in remote sensing. This project assessed its use in the nondestructive evaluation of several blocks of aluminum A356. Hyperspectral data was collected on the blocks after a slight modification to the HIRIS input optics. Several data sets were collected on the specimens and calibration sources. The data sets were averaged to increase the signal-to-noise ratio, processed, and Fourier transformed to produce data




Figure 3. Frame 278 from the raw data cube of the A356 sample used in Figure 2. Note the concentric ring structure of the image formed by the interferograms. The black dots result from pixel imperfections in the focal plane array.





Figure 4. Plot of the logarithm of the magnitude of the eigenvalues of an A356 imagery. There are three types of modes associated with the eigenvalues: modes related to structure of the data (0 to 20); modes associated with system noise (21 to 85); and degenerate modes produced by the processing. (86 to 92).

cubes of spectral data. Preliminary analysis of the spectral data using Principle Component Analysis reveals interesting structure in the data, but additional analysis is required to fully assess the potential of this technique.

Acknowledgments

The author would like to thank J. Bixler for collecting the data, W. Aimonetti for help with the analysis, and C. Bennett for interesting discussions on hyperspectral infrared imagery.

References

1. Bennett, C. L. (1988), "LIFTERS, The Livermore Imaging FTIR Spectrometer," *Fourier Transform Spectroscopy: 11th International Conference*, J. A. de Haseth, ed., American Institute of Physics.
2. Roberts, R. S., W. D. Aimonetti, and J. V. Bixler (1998), "Material Characterization Using a Hyperspectral Infrared Imaging Spectrometer," *Proceedings of the IEE Thirty-Second Conference on Signals, Systems, and Computers*, Pacific Grove, California, November.
3. Beer, R. (1992), *Remote Sensing by Fourier Transform Spectrometry*, John Wiley and Sons, New York, New York.
4. American Society for Metals Reference Book, 2nd ed., 1983.
5. Malinowski, E. (1991), *Factor Analysis in Chemistry*, 2nd ed., John Wiley and Sons, New York, New York.
6. Fukunaga, K. (1990), *Introduction to Statistical Pattern Recognition*, Academic Press, New York, New York. 



n-Situ Identification of Anti-Personnel Mines Using Acoustic Resonant Spectroscopy

Randy S. Roberts and Roger L. Perry

*Defense Sciences Engineering Division
Electronics Engineering*

We have presented a novel technique for discriminating between anti-personnel (AP) mines and other buried objects. The technique is based on measuring the acoustic resonant spectrum of a buried object, using a probe to provide an excitation signal and a stand-off radar to detect the object's response. Several experiments were performed on mine-like objects to assess the potential of the identification technique. One set of experiments found spectral features that might be used as the basis for discrimination algorithms. The second set of experiments indicate that the spectra of the objects have a tolerable degree of variability. Taken together, we conclude that it is highly likely that AP mines can be distinguished from other buried objects using acoustic resonant spectra acquired by a stand-off radar.

Introduction

The detection of buried anti-personnel (AP) mines is a problem of tremendous importance in many parts of the world. The International Committee of the Red Cross and Red Crescent Societies estimates that approximately 110 million AP mines have been placed in 64 countries world wide.¹ The disruption to daily life caused by AP mines in these countries is enormous. For example, over the last fifteen years, land mines in Afghanistan have caused 600,000 casualties (one out of fifty Afghans).

In addition to causing direct injuries, land mines often contribute to malnutrition and promote disease in a population. For example, safe drinking water can be denied to a village where there is a mere suspicion of a land mine field. The village might therefore use unsafe supplies, increasing the risk of dysentery. Similarly, land mines can easily remove farmland from production, increasing the risk of malnutrition.

Modern AP mines are surprisingly simple, effective, and insidious devices. They tend to be cylindrical in shape, with diameters ranging from 6 to 15 cm, and heights from 3 to 6 cm.² They often have simple pressure-sensitive triggers, requiring a mass of 3 to 25 kg for detonation. The explosive charge is typically 50 to 200 gm of TNT. This amount of explosive can cause serious injury to an adult, such as the destruction of a foot or leg, and it can kill small children.

Modern mines contain very little metal, and thus are very difficult to locate with conventional metal detection instruments. Operationally, they are buried in soil to a depth of 1 to 4 cm, and laid out in a variety of patterns, depending on the application. They are very robust devices, and can remain operational for decades after planting.

Conventional demining techniques are time-consuming, labor-intensive and dangerous.² Due to its thoroughness and effectiveness against minimum-metal mines, manual probing is the most widely used demining technique in the world.

For clarification, we note that we draw a distinction between humanitarian demining and mine field clearing. Mine field clearing refers to a military unit breaching a mine field. Often, this is accomplished by a simple linear path through the field. Speed is of the essence. Safety and thoroughness of mine removal are secondary. In contrast, humanitarian demining calls for clearing large areas of terrain, and thoroughness of mine removal is paramount (as is the safety of the demining crew). Speed is a secondary concern.

In manual probing, a deminer pushes a probe into the ground at a shallow angle. The probe is lifted slightly and extracted. A deminer trained in this technique can feel a mine above the probe as it is lifted and extracted. To ensure that all AP mines in an area are detected, probing is performed over a 3-cm-x-3-cm grid. With this type of search pattern, one deminer can clear approximately one square meter of land per day.

Emerging mine detection techniques include infrared sensing for wide area detection, olfactory and chemical sensors for explosive detection, and ground-penetrating radar for detecting buried objects. Of these techniques, ground-penetrating radar has shown the most promise.

A current project at Lawrence Livermore National Laboratory (LLNL) seeks to image buried objects using micropower impulse radar (MIR) as an imaging sensor. The idea is to rapidly survey the top layer of soil for objects that have geometric properties similar to AP mines. When such an object is detected, it is marked for probing. The application of such a ground-imaging system would no doubt increase the efficiency of deminers. However, deminers could easily be inundated by the large number of false alarms produced by shrapnel, debris and other objects that produce images similar to AP mines.

***In Situ* Identification Technique**

The technique described here allows a deminer to discriminate between mines and other buried objects *in situ*. It thus provides a means to distinguish AP mines from other buried objects detected by ground-imaging systems. The technique is based on acoustic resonance spectroscopy (ARS), a technology that has been successfully applied to the identification of chemical munitions.³⁻⁵

The ARS technique is based on the premise that objects of interest, such as AP mines, have characteristic resonant frequencies that we can excite and detect and that are nominally reproducible. Since the resonant frequencies of an object are functions

of its geometry and construction materials, it follows that objects of a similar nature will have similar patterns of resonant frequencies. Thus, the frequency response pattern of an object can be used to identify it.

As mentioned, this technique was used with great success in the ARS Munition Classification System (MCS) developed at Los Alamos National Laboratory. An instrument based on this concept was developed, tested, and accepted for use as a verification tool for the 1997 Chemical Weapons Convention.

There are several ways to estimate the acoustic resonance spectrum of an object. In the case of an AP mine, the objective is to obtain the spectrum in the most non-invasive manner possible. **Figure 1** is a schematic diagram of the technique.

To begin with, an excitation force spanning the frequency band containing the resonances is applied to (or near) the object. The magnitude of the excitation force does not need to be large to obtain a useful spectrum. In the case of the ARS-MCS, the excitation force is provided by a piezoelectric transducer. The induced vibrations on the munition have amplitudes on the order of 10 nm.

For AP mines, a simple approach is to touch the object with a probe that produces low amplitude vibrations over the frequency band of interest. Although this requires contact with the object, it is perhaps the best way to excite the object's resonant frequencies. And as previously noted, touching mines with a probe is accepted practice with deminers.

With regard to the form of excitation signal, a narrow-band sinusoid stepped over the frequency range of interest has a major benefit: the tone provides a coherent reference for processing the received signal. A coherent reference can be used to increase the signal-to-noise ratio of the returned signal, and can also be used to extract phase information.

To collect the frequency response of the object, a sensor is required to detect the object's vibrations. In general, radar is well suited for the non-invasive sensing of vibrations. In particular, the MIR technology invented at LLNL provides a unique means of non-contact, stand-off vibration sensing.⁶

Even though the objects of interest are buried, the burial depths are shallow enough that a low power radar should have little problem sensing the vibrations. By using localization information provided by the LLNL MIR ground-imaging radar, the range gate and other parameters on the vibration




Figure 1. Schematic of the *in-situ* anti-personnel mine identification technique. A probe excites the buried object with a stepped-frequency sinusoidal signal. The response of the object is received by an MIR system, analyzed, and classified by the computer.

sensing radar can be adjusted to maximize the signal-to-noise ratio of the returned signal.

Progress

We have conducted several experiments to assess the feasibility of the identification technique. The experiments consisted of measuring the acoustic resonant spectrum of several objects over a limited frequency range. The measurements were collected with unburied objects to maximize the signal-to-noise ratio.

Four objects of similar geometry were used in the experiments (**Fig. 2**). The first object was an inert M14 mine; the second a nondescript mine detection target (MDT); and the third object, serving as debris, was a galvanized pipe covered with duct tape. A fourth object, not pictured, which also served as debris, was a plastic can containing PlayDough®.

The M14 surrogate is made of a hard plastic, roughly cylindrical in shape and approximately 4.3 cm in height, and 5.2 cm in diameter. The MDT is also made of hard plastic, and is approximately 3.5 cm in height, and has a diameter of 7.5 cm. The pipe is 8 cm in height and has a diameter of 6 cm., and the PlayDough® a cylindrical plastic can approximately 6.1 cm in diameter and 8 cm in height.

The objects were secured in a wooden cradle for the experiments. The excitation signal was provided by an HP35670A signal generator. That signal was amplified and applied to a Wilcoxon F4 vibrator. A probe approximately 1 ft in length was attached to the vibrator, and touched to the sides of the objects. (See **Figs. 3** and **4** for photographs of the experimental arrangement.)

Object vibrations were sensed by a low-power MIR. The radar was placed 60 mm from the surfaces of the objects, in an orientation perpendicular to the surface. Both the drive signal from the HP generator and the return signal from the MIR were digitized at a 40-kHz sampling rate.

Spectra from the objects were collected over the frequency ranges 50 to 100 Hz, 100 to 200 Hz, and 200 to 400 Hz. In each band, the frequency of the drive signal was stepped over approximately thirty tones, each with a duration of 1.5 s. The response signals from the radar were re-sampled to 10 kHz, normalized to the maximum amplitude of the excitation signal, and Fourier transformed.

Magnitude plots of the frequency responses in the 200 to 400 Hz excitation band are shown in **Fig. 5**. The upper plot is the response of the M14 AP mine, the middle plot the response of the MDT, and the

lower plot the response of the pipe. Notice that the three objects produce distinguishable spectra over this frequency range.

The M14 produces a large response in the 325 to 400 Hz region, whereas the responses of the other



Figure 2. Objects used in the preliminary experiments (PlayDough® not shown). On the left is an inert M14 mine; in the center is a nondescript mine detection target; and on the right is a pipe covered in duct tape.




Figure 3. Experimental apparatus used to measure the acoustic resonant spectrum of surrogate mines and debris. The object under inspection is a small pipe covered with duct tape. The MIR system (above the pipe) is typically used for speech analysis.




Figure 4. Another view of the experimental apparatus. The probe is attached to a vibrator that provides an excitation signal to the object.

two objects are much more subdued. In contrast, the 275 to 325 Hz band of the M14 is slightly depressed compared to the same band in the other objects. The responses in the 200 to 275 Hz band are somewhat similar in that they all tend to increase as the frequency increases from 200 to 250 Hz. The mine surrogate and the pipe have similar responses in the 275 to 400 Hz bands with a noticeable spike in the pipe's response at 275 Hz, and null at the same frequency as the mine surrogate.

The other frequency ranges contained similar spectral features. As an aside, spectra from the can of PlayDough® are uninteresting, containing few features. Lack of features in the PlayDough® spectra is not unusual since PlayDough® and its plastic container are non-rigid. Most of the excitation energy is absorbed by the PlayDough®, leaving little response signal.

In addition to the experiments described above, several experiments were performed to assess the variability of the spectral measurements. The experiments consisted of multiple measurements of the spectra of the objects over a 200 to 500 Hz band. The multiple measurements were graphed on a

raster plot for comparison. The spectra showed some degree of variability, but overall the measurements appeared repeatable.

Future Work

Although our preliminary experiments indicated the feasibility of this approach to AP mine identification, several critical issues remain. These issues include the structure of the resonant spectra of AP mines and similar objects in the frequency range above 1 kHz, the variability of the resonant spectra, and collecting spectra from buried objects.

We estimate that the resonant frequencies of the M14 mine and MDT occur at approximately 1.2, 4.8, 10.8, ... , kHz.⁷ (A similar set of resonant frequencies occurs for the pipe and PlayDough.®)

The radar that we used in the first set of experiments had an upper frequency range of approximately 1 kHz. As a result, we were unable to investigate the region above 1 kHz, where potentially unique information about AP mines (that is, their resonant peaks) is located. Concerning spectral variability, robust discrimination algorithms depend




Figure 5. Magnitude vs frequency plots of the three objects in Figure 1: (a) the M14 AP mine, (b) the mine detection target, and (c) the pipe.

on spectral features that can be reliably extracted from resonant spectra. The greater the variability of the resonant spectra, the less reliable the spectral features, resulting in poor discrimination. We collected some data to assess spectral variability, but a large ensemble of spectral data is required for a full assessment.

Finally, although our experiments have laid the groundwork for this identification technique, the measurement of spectra from buried objects remains to be performed. This set of experiments would collect spectra from objects similar to the ones used here, but buried in a variety of soils, such as sand, rock, and clay, and under a variety of conditions, dry to very wet.

We conclude this section with thoughts on the automated pattern classification of spectra and the future of algorithm development.

Automated pattern classification is necessary to provide deminers with confidence in the technique. From the deminer's point of view, the processing should produce a binary decision ("yes, the object is a mine," or "no, the object is not a mine"), along with a numerical measure of confidence.

The key to providing such a decision to the deminer resides in the object's spectral features. If invariant spectral features can be found, clustering techniques can be used for classification. Such was the case in the ARS-MCS liquid-solid discrimination algorithm, where features extracted from the spectra of liquid- and solid-filled munitions grouped into two clusters, one for solids, and one for liquids.

If the spectral features tend to be variable (due to factors such as soil conditions and the age of the mine), a comparative approach can be used. Such was the case in the ARS-MCS munition classification algorithm, where spectral features of different chemical munitions (in this case, the frequencies of resonant peaks) tended to vary. The munition classification algorithm compared spectral features from the unknown munition to templates of spectral features built from known munitions. (The templates quantified the variability of the spectral features.)

While the algorithms developed for the ARS-MCS may not be directly applicable to the proposed project, they form a strong foundation for future algorithm development.

Acknowledgments

The authors would like to thank J. Holzrichter and G. Burnett for the use of their MIR and data acquisition systems, and T. Woehrle for the use of equipment in LLNL's Modal Analysis Laboratory.

References

1. U.S. Army Communications Electronics Command, Night Vision and Electronic Sensors Directorate, Countermine Division, *Humanitarian Demining* (1998). <http://www.demining.brtrc.com>. Also see *Direct and Indirect Consequences of Land Mines on Public Health*, <http://www.demining.brtrc.com/contents.htm>.
2. *Jane's Mines and Mine Clearing* (1997), C. King, ed., Jane's Information Group Inc., Alexandria, Virginia.
3. Roberts, R. S., J. T. Chen, O. A. Vela, and P. S. Lewis (1993), "Munition Classification Using an Acoustic Resonance Spectroscopic Technique," *Proceedings of the Twenty-Seventh Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, California, November 1–3.
4. Roberts, R. S., P. S. Lewis, J. T. Chen, and O. A. Vela (1994), "Techniques for Classifying Acoustic Resonant Spectra," *Proceedings of the Twenty-Eighth Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, California, October 30–November 2.
5. Roberts, R. S., P. S. Lewis, and O. A. Vela (1995), "A Pattern Recognition Algorithm for the Blind Discrimination of Liquid- and Solid-Filled Munitions," *Proceedings of the Twenty-Ninth Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, California, October 29–November 1.
6. Holzrichter, J. F., G. S. Burnett, L. C. Ng, and W. A. Lea (1998), "Speech Articulator Measurements using Low Power EM-wave Sensors," *J. Acous. Soc. Am.* **103** (1) p. 622, January.
7. Skudrzyk, E. (1966), *Simple and Complex Vibratory Systems*, Pennsylvania State University Press, University Park, Pennsylvania. 

An Acoustic Technique for the Non-Invasive *In-Situ* Measurement of Crystal Size and Solution Concentration

Diane J. Chinn and Paul R. Souza

*Manufacturing and Materials Engineering Division
Mechanical Engineering*

Harry F. Robey

Laser Programs

We demonstrated the use of acoustic measurements for tracking potassium dihydrogen phosphate (KDP) crystal growth. Both KDP solution concentration and KDP crystal size can be found by using information derived from acoustic wave propagation in the solution. Acoustic measurements show good correlation to conductivity measurements for KDP solution concentration.

Introduction

The National Ignition Facility (NIF) at Lawrence Livermore National Laboratory (LLNL) is currently developing rapid growth techniques for potassium dihydrogen phosphate (KDP) single crystals. NIF requires over six hundred KDP optical components, sized about 0.4 m², for optical switching and frequency conversion of the beam line. Rapid boule growth rates of approximately 10 to 15 mm/day are necessary to produce the large number of KDP optical components at a reasonable cost. While conventional growth methods would take more than one year to grow a KDP crystal of the size required, NIF scientists are developing a technique to grow the crystal in approximately eight to ten weeks.

Crystals are grown from a point seed in a highly supersaturated solution of KDP. The supersaturation level of the solution is the primary operational parameter used to determine the growth rate of the crystal.¹ Thermal control of KDP solution is used to keep supersaturation at desired levels. Fast growth rates of such large crystals requires stringent controls on the KDP solution concentration.

Several techniques are commonly used to determine solution concentration in crystal growth. The current method uses visual measurement of the KDP crystal size. Given the known solid density of the KDP crystal, the mass of salt in the crystal, removed from solution is estimated. Subtraction of this

estimated value from the initial starting salt gives the solution concentration as a function of time.

This method is rather inaccurate because visual crystal sizing is accurate to only 1 to 2 mm. Other methods using density (from buoyancy), refractive index, and electrical conductivity are also commonly used commercially by crystal growers to determine solution concentration. Generally, solution concentration is determined from empirical relationships of the measured properties. These commercial techniques all require access to the solution. Acoustic measurements offer the possibility of a non-invasive technique for monitoring crystal growth. High-frequency acoustic measurements have the advantage of non-invasively obtaining information about the solution concentration and crystal size during the growth process.

This work assessed acoustic velocity as a diagnostic for KDP crystal growth. Electrical conductivity measurements for solution concentration are used as a comparison. Acoustic velocity measurements are also tested as a means of sizing the crystal during growth.

Progress

We designed a probe to empirically determine the relationship between acoustic velocity and temperature and concentration of KDP solution. Used for determining the empirical relationship, this probe is

designed to operate invasively in the solution. After determining the empirical relationships, acoustic measurements can be made non-invasively from outside the solution.

Figure 1 shows a sketch and a photograph of the probe and the measured acoustic waveform. A 15-MHz ultrasonic transducer mounted on top of the probe sends a wideband pulse (bandwidth = 55 %) propagating into the probe. Fabricated from fused silica, the probe has threaded sides to disperse diffracted waves incident on the sides.

Notched into the probe is a 10-mm gage length. The low coefficient of thermal expansion of fused silica ensures that the gage length remains constant through the temperature range of KDP supersaturation. A pulsed acoustic wave travels roundtrip from the transducer, into the probe, through the solution within the gage length and back to the transducer through the probe. Travel time of the pulse through the KDP solution gives the acoustic velocity. The wideband 15-MHz ultrasonic pulse permits high temporal resolution of the acoustic velocity in the KDP solution. The probe can discern changes as small as 0.25% in acoustic velocity.

In fluids, acoustic velocity is directly related to density. For example, empirical relationships in sea

water show that acoustic velocity increases linearly with salinity and as a second-order polynomial with temperature.² KDP solution should follow acoustical relationships similar to those found in sea water.

Figure 2 shows acoustic velocity and conductivity data for different concentrations of KDP solution. Curve fitting indicates that acoustic velocity is related to temperature through a second-order polynomial. At a given temperature, velocity increases with solution concentration. The data show that acoustic velocity peaks between 60 °C and 65 °C for all concentrations. From linear regression at $T = 60\text{ °C}$,

$$v(C_0) = 0.6414C_0 + 1.534, \quad (1)$$

where C_0 = solution concentration and $v(C_0)$ = acoustic velocity.

For a fixed concentration, conductivity increases linearly with temperature, as expected. At $T = 65\text{ °C}$,

$$c(C_0) = 379.7C_0 + 82.9, \quad (2)$$

where $c(C_0)$ = conductivity.

Empirical equations derived from curve fitting for normalized velocity and concentration are shown in

Figure 1. Probe and waveform. Attached to a 15-MHz transducer, the fused silica probe with 10-mm gage length is placed in the KDP solution to measure acoustic velocity: (a) measurement configuration; (b) measured acoustic signal; and (c) photograph of the probe. Three reflections from the probe are measured in the signal. Velocity is found from the travel time in the KDP solution.

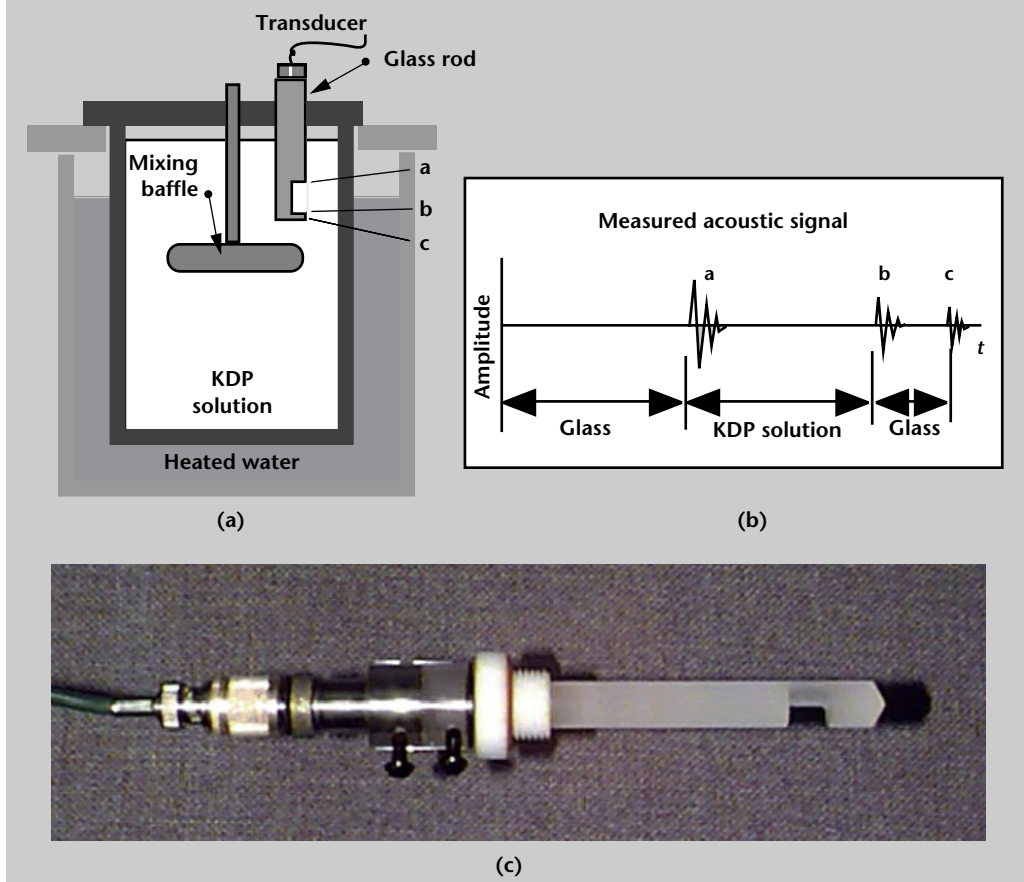


Fig. 3. In **Fig. 3a**, measured velocity, $v_{measured}$ is normalized to $v(C_0)$ by the equation

$$v^*(T) = \frac{v_{measured}}{v(C_0)}, \quad (3)$$

where $v^*(T)$ is normalized velocity. In **Fig. 3b**, conductivity is normalized to $c(C_0)$ by the equation

$$c^*(T) = \frac{c_{measured}}{c(C_0)}, \quad (4)$$

where $c^*(T)$ is normalized conductivity.

Conductivity shows better correlation to concentration ($R = 0.99976$) than velocity ($R = 0.9106$), though for some applications, the non-intrusive nature of the acoustic technique may be of more importance than the small loss in accuracy of the measurement. Also, the conductivity probe is very

sensitive to the presence of bubbles on its measurement surface, which strongly affect the conductivity reading.

The empirical equations shown in **Fig. 3** are useful for *in-situ* continuous monitoring of the solution concentration. Given a temperature, these equations allow us to relate measured values of velocity and conductivity to concentration.

Crystal size in solution is currently measured visually with a telescope and a scale mounted to the side of the tank. The resolution currently obtainable is approximately 1 to 2 mm. Acoustic measurements of crystal size were performed with the configuration sketched in **Fig. 4**. In this configuration, transducers are mounted on the outside of the KDP solution tank. The arrival time of the reflected wave

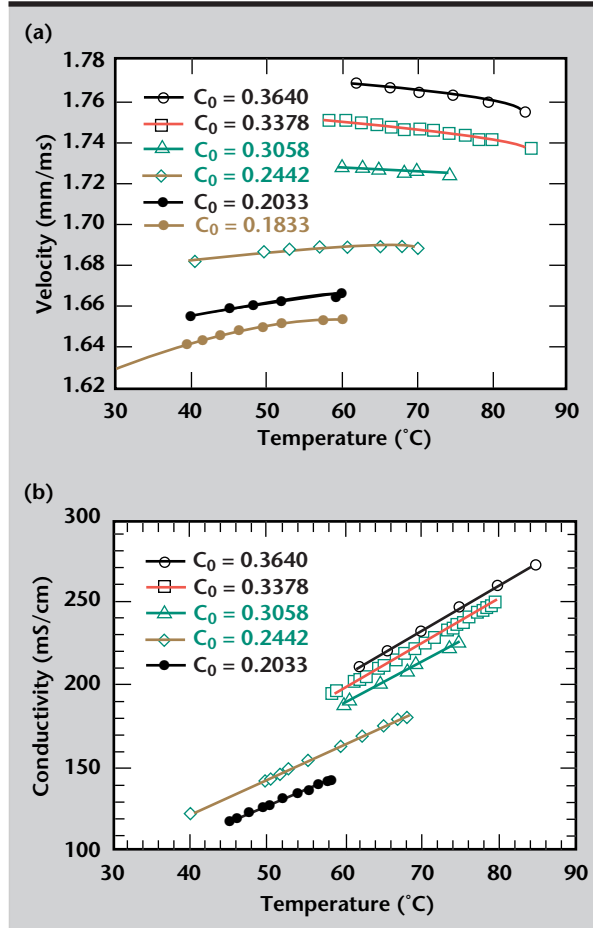


Figure 2. Acoustic velocity (a) and conductivity (b), measured at different temperatures and concentrations. For a given concentration, velocity follows a second-order polynomial fit, while conductivity increases linearly with temperature.

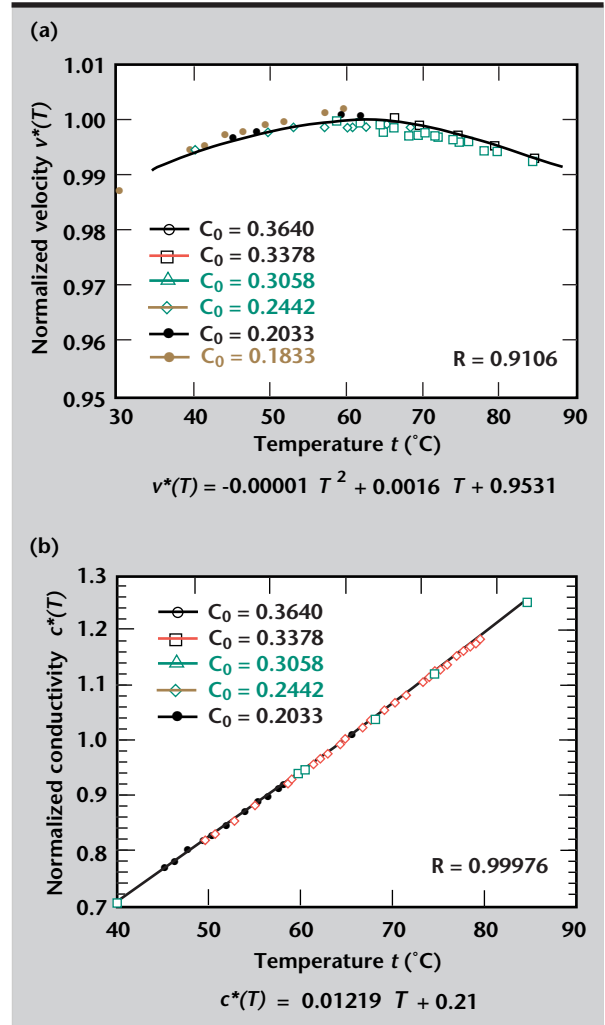


Figure 3. Normalized velocity (a) and normalized conductivity (b) derived from empirical equations. Velocity and concentration are normalized for concentration according to Eqs. 1 and 2. Velocity follows a second-order polynomial fit with temperature with the maximum velocity occurring at 60 °C. Conductivity fits its empirical equation ($R = 0.99976$) better than velocity ($R = 0.9106$).




Figure 4. Illustration of acoustic techniques as a non-invasive diagnostic for crystal growth. Solution concentration is found from a reference wave propagating through the solution. Crystal size is derived from the travel time of the reflected wave from the crystal.

from the crystal face is calibrated to the arrival time of the reference wave. This technique is sensitive to approximately 0.1-mm changes in crystal size, resulting in a factor of 10 improvement.

The configuration in **Fig. 4** shows a non-invasive method of measuring solution concentration in addition to crystal size. Acoustic velocity is obtained from the reference signal. Given the temperature and acoustic velocity, the solution concentration can be determined using the empirical relations in **Fig. 3**.

Summary


Acoustic measurements have proven to be a viable, continuous, in-process diagnostic technique for crystal growth. The technique has no adverse effect on the solution. Acoustic velocity measurements for solution concentration correlate well with electrical conductivity measurements. This work developed a supplemental tool for the continuous measurement of both KDP solution concentration and crystal size.

Future Work

Because of its ultimate simplicity and minimal cost, visual observation is currently the method being used to determine crystal size during growth. Solution concentration, however, is being measured only before and after the crystal growth run. It would be very desirable to have continuous measurement capability to obtain this quantity throughout the growth run.

This work demonstrated the feasibility of using acoustic measurements for determination of the solution concentration. Additional work is needed, however, to minimize the cost of a fielded acoustic measurement system and to provide a more convenient operator interface for converting the raw acoustic signals to a direct reading of solution concentration.

References

1. Cooper, J. F., and M. F. Singleton (1985), "Rapid Growth of Potassium Dihydrogen Phosphate Crystals," Lawrence Livermore National Laboratory, Livermore, California (UCRL-91795).
2. Birks, A. S., and R. E. Green (1991), "Nondestructive Testing Handbook, Ultrasonic Testing," Vol. 7. 

Micro-X-Ray Computed Tomography for PBX Characterization

**Diane J. Chinn, Jerry J. Haskins, Clinton M. Logan,
Dave L. Haupt, and Scott E. Groves**

*Manufacturing and Materials Engineering Division
Mechanical Engineering*

John Kinney

*Materials Science and Technology Division
Chemistry and Materials Science*

Amy Waters

*Department of Material Science
Johns Hopkins University
Baltimore, Maryland*

We compared two micro x-ray computed tomography systems for use in the characterization of plastic-bonded explosive microstructure. This comparison will help guide the development of a technique to generate an accurate structural model of plastic-bonded explosive from nondestructive x-ray tomographic scans.

Introduction

Current efforts to characterize the weapons stockpile have produced heightened interest in the aging mechanisms of plastic-bonded explosives (PBX). Structural modeling offers one method of understanding the aging process. Accurate modeling of PBX crystal and binder components requires a realistic 3-D characterization of the microstructure.

X-ray computed tomography (CT) is an advanced imaging technique that provides 3-D characterization of materials for nondestructive evaluation. *In-situ* characterization of structures and components makes CT an attractive method for analyzing existing parts and components. Recent developments in hardware have permitted CT imaging with micron spatial resolution.

In addition to geometric characterization, CT offers quantitative information on the linear x-ray attenuation of component materials. Related to source energy, material density, and chemical composition, x-ray attenuation can be used to identify and characterize different phases in materials.

Progress

We investigated the use of CT for microstructural characterization of PBX. Two CT systems, KCAT and XTM, were used to quantify the structure of a 2-mm-diameter sample of mock PBX. Component materials of mock PBX do not have the same x-ray attenuation as PBX. However, the microstructure of mock-PBX approximates PBX.

The x-ray tomographic microscope (XTM) at Lawrence Livermore National Laboratory (LLNL) uses beam line 10-2 at the Stanford Synchrotron Radiation Laboratory (SSRL). Located at the Stanford Linear Accelerator Laboratory (SLAC), the XTM facility contains a 15-keV monochromatic synchrotron radiation source tomography system.^{1,2} The XTM system scanned the mock-PBX sample at 5 μm full-width half-maximum (FWHM) of the point spread function spatial resolution.

Recently developed by the Nondestructive Evaluation (NDE) Section at LLNL, the high-resolution KCAT x-ray tomographic system uses a polychromatic source.³ The mock-PBX sample was scanned at an energy of 60 keV peak at 10 μm FWHM spatial resolution.

Figure 1 shows CT slices of the 2-mm-diameter sample from each of the systems. KCAT data is shown in **Fig. 1a**, XTM data is shown in **Fig. 1b**. The slices are representative slices taken at different locations on the sample. Because of higher spatial resolution and higher contrast sensitivity, the XTM slice in **Fig. 1b** provides sharper edges and better contrast of the component structures than the KCAT slice.

Histograms of linear attenuation values for each of the systems are shown in **Figs. 1c** and **1d**. Attenuation values in the XTM data show four distinct phases in the sample. The KCAT data does not distinguish separate phases in the attenuation histogram.

Mock PBX consists of four phases: talc, Kel-F, cyanuric acid, and voids. Expected attenuation values for these component materials at discrete energies ranging from 10 to 100 keV are calculated and presented in **Fig. 2**. For the monochromatic XTM energy of 15 keV, expected attenuation values are 14 cm^{-1} for talc, 12 cm^{-1} for Kel-F, 3.4 cm^{-1} for cyanuric acid, and 0 cm^{-1} for void space. These values are based on 100% density. It should be noted that *in-situ* density of the component materials can deviate substantially from 100% density. Consequently, *in-situ* attenuation values can be different from the expected attenuation values.




Figure 1. KCAT (a) and XTM (b) representative slices, showing the difference in the systems. Histograms of attenuation values in KCAT slice (c) and XTM slice (d) show the difference in contrast sensitivity from the two systems.




Figure 2. Calculated linear attenuation coefficient vs energy for three of the four components of PBX, plotted for comparison to the experimental results. Attenuation is calculated from chemical compositions and density, where Kel-F = $H_2C_8F_{11}C_{13}$ at 2.02, cyanuric acid = $H_{10}C_3N_3O_2$ at 2.50, and talc = $H_2O_{12}Si_4Mg_3$ at 2.78.

In Fig. 3, the four phases of the sample from the XTM scan are segmented from a region of interest (ROI) according to attenuation values. Segmentation reveals the morphology of each of the components. Table 1 shows the range of attenuation values for each of the phases in the ROI. The range of the measured attenuation values in the scan in Table 1 fits well within the range of the expected attenuation values for each of the component materials and allows the identification of each of the phases.

Although the KCAT slice in Fig. 1b shows some features of the microstructure, distinct phases are not readily identifiable in the histogram in Fig. 1c. Unlike the monochromatic XTM source, the polychromatic KCAT source energy ranges from 0 to 60 keV. Figure 2 shows the range of attenuation values for each of the component materials. Contrast sensitivity between the phases is diminished by using a polychromatic source. As a result,




Figure 3. Segmentation of the XTM slice, showing morphological information of the phases. This segmentation is based on one feature, linear attenuation. The original XTM data and all segmented phases (indicated by white areas) are as labeled.

Table 1. Attenuation values in the XTM data.

Phase	Attenuation (cm^{-1})	Slice area (%)	Size (μm)
1. Talc	12.4–22	11.7	7.5–35
2. Kel-F	7–12.4	15.6	10.5–128
3. Cyanuric acid	0.1–7	71.2	—
4. Void	0–0.1	0.7	3.5–30

the phases in the sample are difficult to identify based on their measured attenuation values.

Summary

Micro-CT was assessed for use in the 3-D characterization of PBX microstructure. Comparison of two micro-CT systems indicates that a monochromatic source with 5- μm FWHM spatial resolution provides sufficient spatial resolution and contrast sensitivity to characterize and identify mock-PBX microstructure. Our initial attempt to use a polychromatic source at 10 μm spatial resolution allows imaging of the microstructure but does not easily permit identification of phases in mock-PBX.

Future Work

The XTM micro-CT system is not readily available to LLNL's NDE Section and as such is not an easy option in PBX characterization. This project sought to determine the viability of the KCAT system, a system readily available, for PBX characterization.


Several modifications and processing techniques can be used to improve the sensitivity of the KCAT system to PBX microstructure. The KCAT source spectrum can be modified to a spectrum more appropriate for distinguishing the materials in PBX. This can be accomplished by reducing the peak energy and filtering the source. Spatial resolution can also be improved. A plan for improving spatial resolution of KCAT has been developed and awaits funding.

In addition, image processing techniques can be used to identify phases. Preliminary processing indicates that curve-fitting the attenuation histogram, thereby forcing attenuation values into four phases, may help segment different phases of PBX. Edge enhancement routines can help sharpen the image, improving characterization of the microstructure.

Acknowledgments

Some of this work was performed at SSRL, which is funded in part by the U.S. DOE/Chemical Sciences.

References

1. Kinney, J. H., and M. C. Nichols (1992), "X-ray tomographic microscopy using synchrotron radiation," *Annual Reviews of Materials Science*, **22**, pp. 121–152.
2. Kinney, J. H., D. L. Haupt, M. C. Nichols, T. M. Breunig, G. W. Marshall, and S. J. Marshall (1994), "The x-ray tomographic microscope: 3-D perspectives of evolving microstructure," *Nuclear Instruments and Methods in Physics Research A*, **347**, pp. 480–486.
3. Martz, H., J. Haskins, C. Logan, D. Perkins, D. Rikard, R. Roberts, D. Schneberk, S. Sengupta, and E. Updike (1998), "CT characterization of CVIT-II/III (U)," Lawrence Livermore National Laboratory, Livermore, California (UCRL-ID-130261), June. 



valuation of an Amorphous Selenium Array for Industrial X-Ray Imaging

Clinton M. Logan, Jerry J. Haskins, Kenneth E. Morales, and Earl O. Updike

*Manufacturing and Materials Engineering Division
Mechanical Engineering*

James M. Fugina and Anthony D. Lavietes

*Defense Sciences Engineering Division
Electronics Engineering*

Daniel J. Schneberk

*Computer Applications Science and Engineering Division
Computations*

Gregory J. Schmid

*Isotope Science Division
Chemistry and Materials Science*

Keo Springer

*Applied Research Engineering Division
Mechanical Engineering*

Peter Soltani

*Liberty Technologies, Inc.
Conshohocken, Pennsylvania*

Kenneth Swartz

*Sterling Diagnostic Imaging
Newark, Delaware*

The large market for digital x-ray imaging in medicine has driven the development of flat panel imaging devices. We evaluated one such device using the higher energy x rays typical of industrial applications. This imager is a 360-mm- \times -430-mm array using amorphous selenium (amSe). The defining characteristic of this technology is that electrons produced by radiation interaction in the Se are directly collected and processed with no intermediate scintillation or other conversion process. Comparing performance to existing technology, we found 1) the modulation transfer function (MTF) remains high as spatial frequencies approach the limit imposed by the 139- μ m pixel size; 2) increasing x-ray energy degrades spatial resolution performance somewhat, but the amSe imager performs well in spectra as energetic as 450-kV applied potential; 3) the MTF of this amSe imager drops rapidly with frequency between 0.0 and 0.1 mm⁻¹ for the heavily-filtered spectra used in this study. Secondary radiation transport in material behind the Se is a significant contributor to this attribute; 4) required exposure for the spectra tested ranges from 10 to 100 mR. This is 15 to 100 times less than Lawrence Livermore National Laboratory (LLNL)'s home-built system using scintillating glass and a charge-coupled device (CCD) camera. Cycle time for CT applications will be dominated by read and refresh times of nearly 1 min. This will allow for great flexibility in using microfocus (low power) sources or large source-to-detector distances. Finally, we found 5) except for issues of peripheral electronics, which we did not address, Monte Carlo simulations suggest that this amSe imager could offer attractive performance at x-ray energies of 3 MeV or more.

Introduction

Sterling Diagnostic Imaging (Newark, Delaware) is introducing a new type of x-ray imager for medical imaging. It is designed for use with medical x-ray spectra of about 80-keV average energy. It has an active area ($360 \text{ mm} \times 430 \text{ mm}$) equal to standard chest film, with superior speed and spatial resolution. It is packaged into a box not much larger than a film cassette, and packs an impressive 7.5×10^6 pixels with 14-bit data depth.

This new imager uses manufacturing technology similar to that used for flat panel displays. There is a very thin coating of electrical conductor on the radiation side of the panel, followed by $500 \text{ }\mu\text{m}$ of amorphous selenium (amSe). Beneath each pixel is a capacitor to store charge and a thin-film transistor for read-out. A voltage of a few kV is applied across the amSe so that when x rays produce ionization in the amSe, the resulting charge moves directly along electric field lines and is collected in the pixel in which the radiation interaction occurs. This is shown schematically in **Fig. 1**.

This direct collection of radiation-induced charge is the most important attribute of this new imager. It converts x rays to electrons, which are then collected and integrated without intermediate conversion to visible light. This eliminates the blurring that inevitably occurs in scintillator-based imagers. The technology is scalable to still higher spatial resolution.

The workhorse system configuration for digital x-ray imaging and computed tomography at LLNL is a scintillating glass that is mirror/lens coupled to a CCD camera. Compared to this technology, we expect amSe to be more portable and to have 1) higher x-ray sensitivity, 2) larger detector area, 3) more pixels, 4) higher spatial resolution (expressed relative to the field of view), and 5) fewer alignment issues.

There are several unknowns with respect to the application of amSe to nondestructive evaluation

(NDE) applications. AmSe has only been demonstrated at relatively low x-ray energy used for medical imaging. Most NDE applications use higher-energy x-ray spectra. More energetic x rays will penetrate deeper into and through the panel substrate to the electronic components behind the substrate with unknown effects. Also, resolution may degrade at higher x-ray energy as a result of secondary radiation (x rays and Compton electrons) transport.

In addition, dark current in the imager may be an issue with the longer integration times required for industrial imaging. Finally, x rays must be turned off during read-out. We need to understand the issues associated with computer control of on/off with respect to our x-ray sources.

Progress


We evaluated 1) spatial resolution, 2) dark current, 3) reproducibility, 4) latent image effects, and 5) dynamic range. In addition, we used the Monte Carlo code GEANT to model radiation transport in the detector up to MeV x-ray energies relevant to imaging high opacity objects such as turbine blades, nuclear weapons, and engine components. This work is intended to suggest what the spatial resolution performance might be at MeV energies and indicate possible solutions to shielding the electronics.

We evaluated possible methods for turning x rays on and off for our 450-kV Philips tube and for the 9-MV electron linac.

X-ray Spectra

We acquired data using three different x-ray spectra (**Fig. 2**), filtering to narrow the energy range. The softest x-ray spectrum tested used a tube potential of 120 kV and $76 \text{ }\mu\text{m}$ of W filtration; the intermediate a tube potential of 250 kV and 3.2 mm of Cu filtration; the hardest a tube potential of 450 kV and 6.5 mm of Cu filtration.

Figure 1. Schematic representation of the layer structure and detection method of an amSe imager.



Source Stability

The amSe imager requires that x rays be absent except during the manufacturer-set interval during which the image is acquired. The only method available to us for these tests was to use the preset exposure-time-mode of the LLNL Philips model 451 x-ray source controller. This method of control causes the high voltage to be ramped up and down to turn x rays on and off. To observe drift or change in x-ray sensitivity of the amSe imager, it was necessary to establish that exposure controlled in this way is sufficiently reproducible. Details of this evaluation are given in Reference 1.

To evaluate the reproducibility of exposure controlled from the tube controller, we ran a series of exposures for both the lowest and the highest kV spectra used for the amSe detector tests. The data indicated sufficient source stability to do the detector characterization without an independent measurement of source output for each image.

Monte Carlo Modeling Results

We used the Monte Carlo code GEANT² to model radiation transport in the detector. We chose to model monoenergetic x rays at energies of 70, 120, 170, 450, 3000 and 9000 keV. The first three energies were selected to approximate the effective energies of the polyenergetic spectra that we tested experimentally; 450 keV is the maximum x-ray energy produced in the hardest spectrum we tested (**Fig. 2**); we chose 3000 keV because it is a technologically important energy for large dense objects; and 9000 keV is the end-point of the spectrum produced by our 9-MV linac.

We calculated several configurations to help us understand how the imager might be improved for more energetic x-rays. Results from two configurations are presented here. The first configuration is the array as it exists, while the second explores the effect of adding a Au layer in front of the Se.

The calculated interaction probability for 0.5 mm of Se as a function of photon energy is given in **Fig. 3**. By putting material in front of the Se, it is possible to suppress the response to low energy x rays while possibly increasing the response to high energy x rays. The potential increase at high energy comes from transport of secondary radiation (primarily electrons) out of the Au and into the Se. To have the desired effect, the material should be dense and high-Z. We used computational Au since it is free.

For these calculations, a monoenergetic x-ray source was made incident normal to the detector

surface along a line. The structure of the amSe imager was represented by 12 layers of various materials. The Se is 0.5 mm thick in the Sterling Diagnostic Imaging design. GEANT transports secondary photons and electrons. We present here the results of energy deposition in the Se. The problem was zoned so that we could tally energy deposition as a function of distance from the source. We used 10- μ m-thick lateral spacing and extended the calculation to a distance of 10 mm from the source. To avoid edge effects, we tallied energy deposition in a central 3-mm band within a 1-cm-wide panel. A schematic of the calculation geometry is shown in **Fig. 4**.




Figure 2. Calculated x-ray spectra used to evaluate and compare imaging systems. Spectra are in units of photons/eV-s into each pixel for the source-to-detector distance and tube current used with the amSe imager. The 120kV, 250kV, and 450kV spectra are integrated over 11, 25, and 37 d photons/s. pixel, respectively.




Figure 3. Calculated probability of x ray interacting in a layer of Se 0.5 mm thick.




Figure 4. Schematic of configuration used for Monte Carlo calculations. Tally zones are within a 3-mm band centered in a 10-mm band.




Figure 5. Calculated energy deposition in Se layer of amSe imager as a function of incident x-ray energy.




Figure 6. Calculated energy deposition in Se layer of amSe imager as a function of distance from line source. Zones in the calculation are 0.01 mm wide.

Figure 5 presents the calculated energy deposition for the two configurations reported and the six different x-ray energies. Energy deposition by primary and secondary radiation produces electrons in the Se that are then collected by the applied voltage. There may be some variation in recombination with the ionization density of the secondary radiation, but, to first order, the signal in the detector should be proportional to the energy deposited by the incident radiation. The calculations indicate that in the present configuration, the amSe panel should have a decreasing response (per incident x ray) over the energy range that we tested experimentally.

The location of the minimum at 450-keV x-ray energy is highly approximate since we calculated a sparse energy set.

Addition of the Au gives part of the desired effect on energy deposit. It suppresses the amSe imager response at 70 keV by two orders of magnitude. The response at 450 and 3000 keV is nearly unchanged. For imaging high-Z objects with the LLNL 9-MV linac, this may result in substantially better imaging performance. Most of the important imaging photons are between 1 and 4 MeV. Most of the scatter is below 450 keV. It is disappointing that the response at 3 MeV is not improved, but we tried only one material and thickness. The important conclusion is that a tool exists to alter the energy response of this imager.

There is, however, a potential advantage of the altered configuration for imaging 3-MeV x rays. For the amSe imager in the present configuration, it is possible that the energy deposit comes from a small number of high-energy events. This could make the image quantum-noise limited. The configuration with the Au has much higher interaction probability, possibly leading to more interaction events being recorded and therefore to lessened quantum noise. Insight into these effects is available in the Monte Carlo calculations already performed. Time and budget did not permit this evaluation.

The shape of the energy deposition as distance from the (line) source increases is given in Fig. 6 for the lower energies computed. The zone width in the Monte Carlo calculation was 0.01 mm.

Figure 7 shows the shapes of the energy deposition binned into pixels the size of the amSe imager. The data have been normalized so that each energy has the same integral. These data suggest that there should be an experimentally observable effect on the line-spread function from secondary radiation transport as x-ray energy is increased up to a tube voltage of 450 kV. They also show that the line-spread function will be further broadened by secondary

radiation transport at energies produced by LLNL's 9-MV linac.

Figures 6 and 7 do not reveal one very important aspect of the computational results. A significant portion of the total energy deposited is contained within the long tails of the line-spread functions.

Figure 8 presents the computational results for the configuration of the commercial amSe imager as running integrals. The ordinate is the fraction of the total deposited energy that is contained within a distance, x , of the line source. We have normalized these curves to a value of 1.0 at 10 mm from the source.

In the type of display given in the figure, the desired response is an immediate rise to 1.0 at a very small distance, and then no additional energy deposited at larger distances from the incident radiation. This spread represents only the effects of secondary radiation transport in the amSe imager. It does not include object scatter, air scatter, electronic crosstalk, source unsharpness, digital pixel sampling and many other things.

We see from **Fig. 8** that significant energy is deposited many pixels from the location of the incident radiation. For example, 10% to 30% of the energy is deposited more than 1 mm from the incident radiation, even at medically relevant energies. The importance of this can most easily be appreciated if one imagines imaging a small, high-opacity, object within an otherwise empty field of view.

Even if no source photons reach the amSe imager in the shadow of the object, the aggregate effect of inscatter from the other 7.5 million pixels will cause significant detector response in the shadow. This can severely limit the useful dynamic

range and contrast sensitivity in NDE applications. This same characteristic exists to some extent in every technology available for quantitative digital x-ray imaging.

Results for the configuration with 1 mm of Au inserted before the Se are shown in **Fig. 9**. The curves for lower energies are degraded from the commercial amSe configuration. This is of only intellectual interest since the response is so strongly suppressed. The two most technologically important cases are 450 keV and 3 MeV. These are compared in **Fig. 10**. Clearly the presence of a Au converter layer in front of the Se will have negligible effect on spatial resolution at 450-keV and 3-MeV x-ray energy.

These Monte Carlo calculations show that secondary radiation transport in the imager itself creates long tails on the line-spread function of the amSe imager in its commercial configuration for all energies simulated, 70 keV to 9 MeV.

For imaging with the LLNL 9-MV linac, spatial resolution will be degraded compared to lower energies, but potentially better than any other available technology for area imaging at these energies. Adding a high-Z converter in front of the Se is a potentially viable method of adjusting the imager x-ray energy response with negligible impact on spatial performance.

amSe Image Acquisition Procedure

Initially we had intended to collimate the x-ray beam to approximately one quadrant of the amSe imager. In the event that our tests did irreparable damage, this strategy would have retained some




Figure 7. Calculated energy deposition in Se layer of amSe imager as a function of distance from line source. Data presented here are binned to the pixel pitch of the amSe imager.




Figure 8. Calculated shapes of the integrated energy deposition as a function of distance from the line source for the amSe imager in the commercial configuration.

functionality. However, the imager has an internal calibration procedure that must be run for each x-ray spectrum. To complete this calibration it is necessary to fully illuminate the entire panel. We removed all collimation from the source and allowed the x rays to fill the entire panel and surroundings.

To provide shielding for peripheral electronics within the imager, we applied a 0.25-in. Pb shielding in front of the wide “borders” of the imager. We have insufficient knowledge of internal construction to know whether this shielding protected the vulnerable components. This Pb was installed after the data

set at 120 kV had been acquired, but this is of no consequence, since 120 kV is within the normal operating envelop of the amSe imager.

Spectral and exposure data is presented in **Table 1**. Prior to acquiring the set of images for each spectrum, we performed the “calibration” procedure as outlined by the manufacturer.

We see from **Table 1** that exposure times are quite short. If we were to cut the distance to 1.5 m and run at maximum current, exposure times would be well under 1.0 s. If we operated without source filtration, it would approach 0.1 s. To make full use of this amSe imager in the




Figure 9. Calculated shapes of the integrated energy deposition as a function of distance from the line source for the amSe imager with 1 mm of Au inserted in front of the Se layer.




Figure 10. Effect on energy deposit shape from adding 1 mm of Au in front of the Se.

Table 1. Spectral and exposure data for amSe imager.

Tube voltage (kV)	Filter material and thickness (mm)	Tube current (mA)	Exposure time (s)	Source-to-detector distance (m)	Exposure (mR)
120	W: 0.076	1.75	10.0	3.00	10.2
250	Cu: 3.2	3.00	10.0	3.00	40.8
450	Cu: 6.5	2.00	12.0	3.00	102

Table 2. Spectral and exposure data for PCAT system.

Tube voltage (kV)	Filter material and thickness (mm)	Tube current (mA)	Exposure time (s)	Source-to-detector distance (m)	Exposure (mR)
120	W: 0.076	7.50	110	1.52	1873
250	Cu: 3.2	3.50	100	1.52	1854
450	Cu: 6.5	2.00	85	1.52	2814

future, it is imperative to have some type of shutter control over exposure. These exposure times are less than the time required for the tube to reach operating voltage if we use the tube controller as the on/off switch.

Glass CCD Test Procedure

For the comparative images acquired with the PCAT system, we set up PCAT so that the pixel size was 103 μm , slightly smaller than the pixels in the amSe imager. This resulted in a 105-mm field of view. We used the fiber optic scintillator that is 140-mm square and 12-mm thick. Parameters are given in **Table 2**.

The objects we imaged are 1) a line-pair gage; 2) a Ta cylinder; 3) a Ta edge; and 4) a 3/8-in. socket head screw with nut (amSe imager only).

The overall configuration of the set-up for data acquisition is shown in **Fig. 11**. The amSe imager with objects in place is shown in **Fig. 12**.

With PCAT we omitted the bolt. We imaged a Pb ruler to establish the pixel size and field of view. The Ta edge and the cylinder were imaged together. The line-pair gage was imaged separately.



Figure 11. The overall arrangement used for data acquisition with the amSe imager. The Philips tube head with high-voltage cables is on the left. Three meters away on the table is the imager.

Experimental Results

Exposure Results. We see from the exposure conditions in **Tables 1** and **2** that the amSe imager reaches the desired digital signal level with considerably less exposure than the scintillating glass/CCD system of PCAT. We can correct the PCAT exposure values to account for the difference in pixel size between the two systems. For optically coupled systems such as PCAT, the exposure required is inversely proportional to the pixel area (constant number of x rays per pixel). The exposure advantage derived from **Tables 1** and **2**, corrected for pixel area, is given in **Table 3**.

Dark Image Results. To evaluate the magnitude and stability of the dark image from the amSe imager, we manually selected a region 300×600 pixels that contained few bad pixels. These images were acquired with the same integration time used for x-ray images. The mean and standard deviation for the selected region are given in **Table 4**. The last two digits of the image indicate the minute of image acquisition.

Two of the dark images, 1607 and 1609, have significant counts, about 1% and 2% of the usable dynamic range of the imager. We don't understand




Figure 12. Close up of the amSe imager with objects in place. The white foam was used to support some objects. The white marks on the imager face denote the active imaging area of the device.

Table 3. Exposure advantage, corrected for pixel area.

Tube voltage (kV)	Filter material and thickness (mm)	AmSe exposure advantage (X)
120	W: 0.076	100
250	Cu: 3.2	25
450	Cu: 6.5	15

Table 4. Dark image results.

Image	Mean	Standard deviation
1605	0.00374	1.38
1607	70.5	11.6
1609	171	17.4
1612	0	0
1614	0.0001	.042
1615	0.0003	0.011

this behavior. We believe that the images written to disk have been dark-corrected by the Sterling software, so we expect an average of zero, with a variance caused by noise in the dark images.

The images appear to have been clipped so that no negative values appear. We also cannot rule out the remote possibility that images 1607 and 1609 may have been affected by external interference. Four of these images are effectively dead zero with no significant standard deviation. This needs more work.

Temporal/Radiation Stability. One of the problems in our scintillator-based systems is complex temporal behavior. The particular glass formulation that we use has both a long-lived “afterglow” and light output that increases with recent radiation exposure. We checked the amSe imager for a residual image within a dark image after a series of five

object images. Some residual image is observable. This is shown in **Fig. 13**.

The object image at the top left is the fifth in a series of five acquired using the 120-kV spectrum. A portion of the line-pair gage is shown along with a line-out in units of digital levels. The lower image is a dark image acquired as soon as our manual procedure allowed (~2 min) after the fifth in the object-image sequence. The two images are shown with different gray scales to reveal the residual image. The magnitude of the residual image is quite small, within the variation in overall dark image variation in **Table 4**. This effect disappears entirely in a second dark image.

Another form of temporal bad behavior is for the imager signal to exhibit dependence on prior exposure. To evaluate this, we acquired five blank images for each x-ray spectrum tested. These were acquired as rapidly as our clunky manual method and the amSe imager control software would allow, typically one exposure every 2 min. We examined a 300- \times -600-pixel patch from each image and determined mean, standard deviation, median and number of “bad” pixels, where “bad” is defined as having a digital level >15,700 or <2000.

Mean digital values are plotted in **Fig. 14** for three sequences of five blank images, one sequence for each x-ray spectrum. These variations are very large, greater than 10% for the

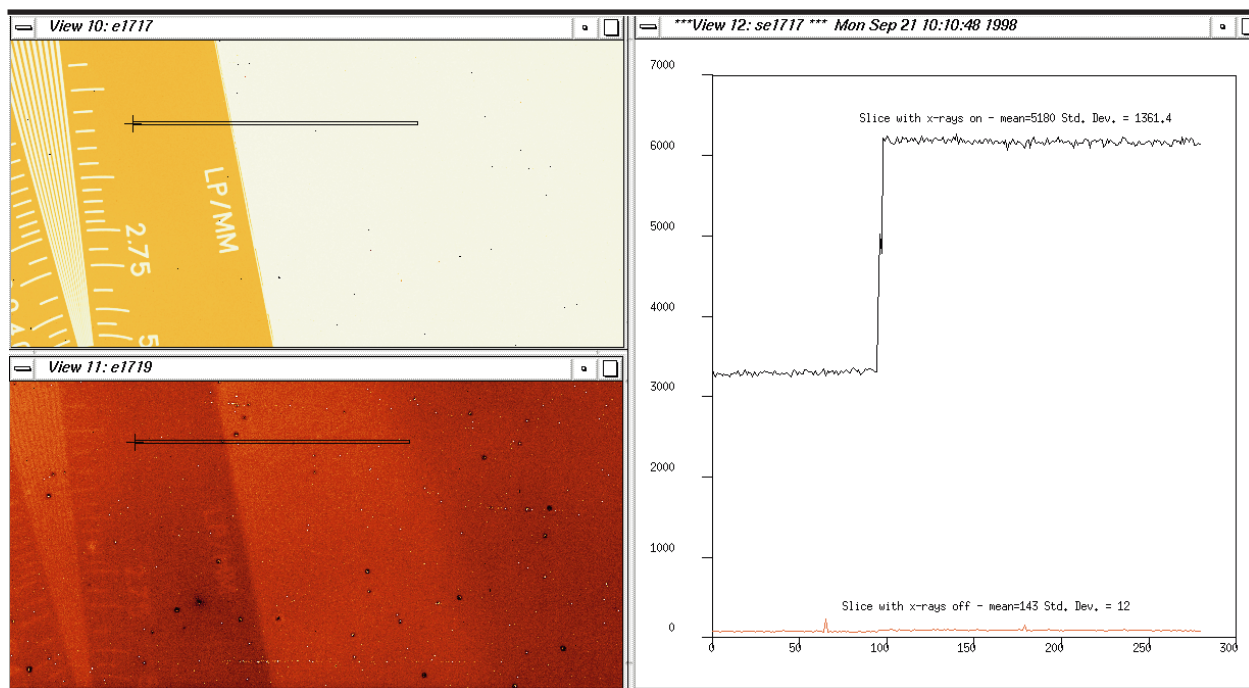


Figure 13. Upper left image is the fifth in a series of five images acquired at 120 kV. The object is a 0.1-mm-thick Pb line-pair gage. Just after acquiring image 5, the dark image, lower left, shows a slight residual image. Line-outs are shown on the right for both images. Line outs were taken at exactly the same pixels in the two images.




Figure 14. Mean digital value for a series of five blank images taken for each of the three x-ray spectra used in these experiments. Note the plot has non-zero origin.

120-kV spectrum. Nothing about the time or sequence behavior makes sense to the authors. The sequence at 120 kV starts high and drops throughout the sequence. The others drop, then climb. The number of bad pixels is very small, 0.03% in the worst case. No trends with continued exposure are evident. The source output is reproducible to about 1%, so the variation of Fig. 14 must arise in the amSe imager. If this is a change in sensitivity that is spatially invariant, then correction is a simple matter. LLNL software routinely renormalizes CT data. Further work is needed to assess whether this variation has any spatial dependence, which would be a serious problem for CT.

Modulation Transfer Function. The MTF is a frequency-domain description of the spatial resolution of an imaging system or component.³

One derives MTF from the image of an edge. First, a line-out is taken. The derivative of the line-out is the line-spread function. The Fourier transform of the line-spread function is the MTF. One consequence of the digital spatial sampling in a digital system is that aliasing can occur if the image contains frequencies that exceed one-half the sampling frequency. For the amSe imager with 139 μm pixels, this limiting frequency is 3.6 mm^{-1} , called the Nyquist limit.

Experimentally determined MTFs for the amSe imager are given in Fig. 15 for all three spectra. They drop very rapidly at low spatial frequency, then hold up well as the Nyquist limit is approached. The MTFs decrease with increasing x-ray spectral energy.

Our Monte Carlo modeling simulates secondary radiation transport, one of the components of the experimental MTF of Fig. 15. Recall




Figure 15. MTFs for three x-ray spectra for the amSe imager. These MTFs were derived from the images of a Ta edge.

that we modeled monoenergetic x rays of 70, 120 and 170 keV as representative of the three polychromatic spectra. On the basis of relative energy deposition at large distances from the incident x ray, the Monte Carlo results indicate that MTF for the 120-kV spectrum will drop less rapidly at low spatial frequency than will the MTFs for the two harder spectra. We see this behavior in Fig. 15.

Experimental MTFs for the glass/CCD imager at the three x-ray spectra are given in Fig. 16. The effect of harder x-ray spectra is in the same direction as for the amSe imager, though of smaller magnitude. Comparing Figs. 15 and 16, we see that amSe is inferior at low spatial frequencies and superior at high spatial frequencies.




Figure 16. MTFs for three x-ray spectra for the glass/CCD imager. These MTFs were derived from the images of a Ta edge.

cies. This comparison is shown in **Fig. 17** for the 120-kV spectrum.

The MTF from digital spatial sampling is given by the expression:⁴ $MTF_x = \sin x/x$. In this expression $x = \text{spatial frequency (mm}^{-1}) \cdot \text{pixel size (mm)} \cdot \pi$. We calculate this quantity for 139- μm pixel size and present it along with the measured amSe MTF in **Fig. 18**.




Figure 17. Comparison of MTFs for the two imaging systems using the 120-kV spectra.

MTF at Low Spatial Frequency. The MTF of the amSe imager at low spatial frequency generally decreases contrast and messes up quantitative measurement. This is made clear by our measurements of signal level in the shadow of an opaque cylinder. The signal at the center of the cylinder arises when the MTF of the imager at low spatial frequencies is not quite unity. This

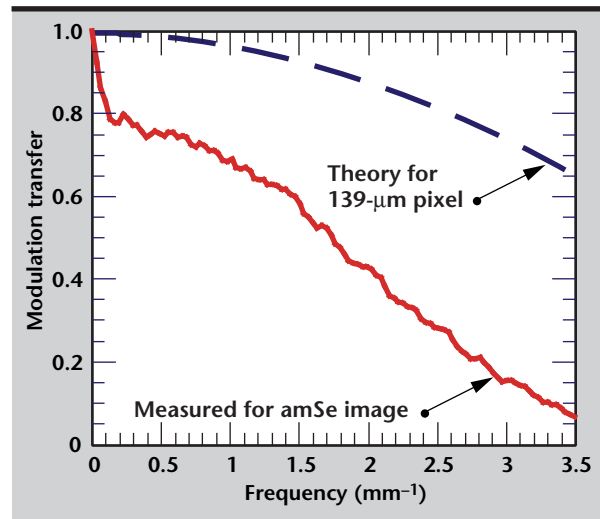



Figure 18. Comparison of measured amSe imager MTF at 120-kV spectrum and theory, for a digital sample size of 139 μm .

Figure 19. Normalized line-outs of the image of the thick Ta cylinder, with both amSe imager and the LLNL PCAT system using a fiber optic scintillating glass lens-coupled to a CCD. The data is split at 15 mm with the amSe imager on the left and the glass/CCD on the right. The digital counts have been normalized so that the maximum level is the same in each signal. Counts for glass/CCD system have had the dark image subtracted. Counts for the amSe imager have internal dark image subtraction.



attribute is important to NDE applications, especially to CT.

The results of these images for the amSe imager are compared to the LLNL glass/CCD system in **Fig. 19** containing portions of line outs from six images (two imagers and three spectra) of the thick cylinder. Each signal has been normalized so that they all have the same maximum value. The trend with increasing spectral energy for the amSe imager is consistent with the Monte Carlo results presented earlier.

Example Images. Images of the Pb line-pair gage taken with the two imaging systems are presented in **Fig. 20**. These images were acquired at 450-kV applied potential. Since the Pb thickness in the gage is only 0.1 mm, this is a modest-contrast object.

Source Control

The amSe imager is continuously active and therefore will not operate as a snapshot camera device in a constant x-ray field. X-Ray exposure time must be controlled with the source or a shutter, and x rays must remain off during read

out, image transfer, and preparation of the imager for the next exposure. To imagine future use of this imager for CT, this whole process must be controlled from the computer used for data acquisition and motion control.

One component of this project was to understand issues regarding source control so that we could lay out a reasonable plan for future use of the amSe imager as the data acquisition component of an NDE CT system.

Though our experiments were not extensive, it appears that the LLNL linacs can be operated in a mode amenable to computer control of relatively short bursts. The ramp-up and ramp-down times appear to be immediate, though a closer examination may find subtle differences.

We performed these tests using a pulse burst that produced about 1 R. At the time these tests were performed, we did not have amSe exposure results for softer x-ray spectra nor the Monte Carlo calculations. These combined results suggest that a 2-s linac burst at 100 Hz may be too long. More detailed testing will be required to pin this down and to evaluate operation at lower




Figure 20. Images of the Pb line-pair gage in the region of 2.0 lp/mm. (a) Acquired with the LLNL PCAT system, using a fiber optic scintillator that is mirror/lens-coupled to a CCD. The only correction applied to this image is dark image subtraction. The brightness gradient from left to right, and the pattern noise are of no consequence because they will disappear once the image is ratioed to the blank image. **(b)** The same region of the line-pair gage, acquired with the amSe imager. The two images are different size because they were acquired with different pixel size and we chose to print them with constant pixel dimensions.

Table 5. Summary of performance of CCD-based configuration.

Feature	amSe	CCD	Comment
Spatial resolution (MTF at 1.0 mm ⁻¹)	0.40	0.15	
Maximum field of view (cm)	43	30	a
1.5-m image cycle time (s)	46	120	b
3.0-m image cycle time (s)	64	375	c
24.7-mm contrast ratio	27	80	d

a) amSe provides spatial resolution and field of view simultaneously. CCD systems can achieve 30 cm field of view, but not simultaneously, with 0.15 MTF at 1 mm⁻¹. LLNL maximum field of view for CCD system is 15 cm.

b) Sum of exposure and read/refresh time. For Philips 451 source operated at 2.0 mA, 450 kV, with 6.5 mm Cu filtration. Source-to-detector distance = 1.5 m.

c) Same as (b), except source-to-detector distance = 3.0 m.

d) Contrast ratio = (digital level in open field ÷ digital level at center of 24.7-mm-diameter opaque disk).

frequency and/or shorter pulse burst, but it looks promising.

Conclusion

This amSe imager shows great promise for digital radiography and CT. We encountered no problems in operating it in x-ray spectra up to 450-kV applied voltage. **Table 5** provides a highly simplified summary of performance relative to one LLNL CCD-based configuration at the 450-kV spectrum tested in this work.

Future Work

The next steps toward routine use of the amSe imager for LLNL Programs are:

1. Write LLNL software to control the imager.
2. Understand the dark image anomalies that we observed.
3. Determine whether the variation we see in blank images is spatially dependent.
4. Test the 450-kV shutter we designed for reproducibility of x-ray illumination.
5. Understand the impact of the poor MTF at low spatial frequency and what mitigation measures can be applied.


6. Perform the same suite of characterization measurements at the 9-MV linac.

7. Evaluate other LLNL CCD-based configurations and commercial detectors based on amorphous Si.

Acknowledgments

J. Mahler (LLNL) made this project possible by catching a bit of our enthusiasm for this technology.

References

1. Logan, C. M., "Evaluation of an Amorphous Selenium Array for Industrial X-Ray Imaging," in press.
2. Brun, R., F. Bruyan, and M. Maire (1987), "GEANT3 Users Guide," DD/EE/84-1, CERN.
3. Bushberg, J. T., J. A. Seibert, E. M. Leidholdt, Jr., and J. M. Boone (1994), "The Essential Physics of Medical Imaging," Williams and Wilkins.
4. Barrett, H. H., and W. Swindell (1981), "Radiological Imaging," Academic Press, New York, New York. 



LANDMARC Radar Mine Detection

Stephen G. Azevedo, Jeffrey E. Mast, and James M. Brase

*Laser Engineering Division
Electronics Engineering*

E. Tom Rosenbury

*Defense Sciences Engineering Division
Electronics Engineering*

This report describes work done in the second year of a three-year project, LANDMARC, for detection of land mines using advanced radar technology. The primary objective of the overall project is to study the use of small radar sensors, such as Micropower Impulse Radar (MIR), and complementary sensors, with optimized detection efficiency. Besides detector design and experimental tests, other aspects of radar mine detection are under study, including electromagnetic modeling, signal processing, assisted detection, and parallel processing methodologies. Both military and humanitarian demining goals were set as standards of operation so that, in the end, a feasible mine detection system could be designed that would meet the goals and constraints of deminers. A first-prototype system was also used to measure detector performance in blind tests. The MIR mine detector has been shown to effectively detect and locate small plastic anti-personnel land mines (with explosive simulant and no detonators) in three different soil types and under adverse conditions. It achieved an 85% overall probability of detection in 18 blind tests, compared with 44% with an Army-standard metal detector, and roughly half the false alarm rate.

Introduction

Land-mine detection is important for soldiers facing an unknown threat, as well as for civilians returning to their land after a conflict. Besides the military threat for our troops overseas, land mines throughout the world pose an enormous hindrance to economic and environmental stability, and indirectly to national security. Policy decisions regarding the development, stockpiling, use, and clean-up by warring nations are being carried out through high-level international negotiations, but there are many problems that still remain—particularly the millions of mines left in the ground. As a national laboratory developing science and technologies in the national interest, Lawrence Livermore National Laboratory (LLNL) has undertaken several efforts over the years to address the land-mine problem.¹⁻¹⁰

The main problem is that roughly one billion mines are still in place, killing large numbers of civilians and children daily. They will continue to be a problem until they can be easily found and removed. Detection and clearance technologies are under investigation by

several agencies, but none are adequate for all situations, particularly for humanitarian demining scenarios. This report describes internal research and development for a new mine detection strategy that takes advantage of the Micropower Impulse Radar (MIR) developed at LLNL.

We describe the second year of this project and the plans to continue toward its deployment for use around the world. We begin with an introduction to the mine detection problem and alternate approaches.

Mine Detection Problem

The mine detection problem being addressed in this project is repeatedly listed “number one” by people who clear minefields: finding the small plastic mines near the surface to be removed after a regional conflict. Other problems are difficult as well (for example, locating a field containing mines, education and care of locals, clearing vegetation, disposing of the mine once found, and verification that a field is cleared), but detecting individual mines is paramount.

This problem represents one of the major difficulties facing humanitarian demining operations, whose goals are to clear land and save lives. It is also a problem for U. S. forces conducting peace-keeping operations (PKO) in troubled areas around the world. The current methods to find these mines are unreliable and slow.

Some facts about this problem are given below.

1. *Anti-personnel (AP) mines are small, have low metal content (a few grams), and are the most numerous.* These mines operate indiscriminately, killing or maiming 2000 civilians and children per month worldwide—draining the resources not only of their own countries, but also of relief efforts and eventually the global economy. One of the most important needs of “deminers” (people who remove post-conflict land mines around the world) is for improved detection of individual minimum-metal AP mines. These mines are numerous, small, effective, and nearly invisible to most current technology. Large anti-tank mines are detectable, and therefore not a problem.
2. *Metal detectors and mine probes are used in 100% of clearance operations.* These tools are primitive, and cause numerous false detections, thereby slowing the operation. In one instance in Cambodia, deminers detected 70 mines in one acre of land, but they also found 500,000 metal fragments in the same field. This false alarm rate slows operations and causes fatigue that can lead to accidents.
3. *Radar technology has never been used by deminers.* Even though many disjointed efforts have tried to develop radar-based mine detection systems in the last 20 years, none are in use. This is because radar has been large, heavy, expensive, requires power, and has been susceptible to clutter. The wavelengths used have been much larger than the AP mines, thus mines have been below the diffraction limit.
4. *Clutter is the limiting factor in all mine detection technologies.* The sources of this clutter are many – surface effects, mine-like targets, electronic or thermal noise, and natural features such as roots, rocks, and holes. The clear need is for improved methods of reducing the effects of clutter for the technologies of interest.

Approaches

Many people in the international humanitarian demining field agree that the most expedient way to address the mine detection problem is to develop a

hand-held, cheap, rugged, and easy-to-use sub-surface mine detector. To do this involves development of better detection strategies.

Current technologies are either incomplete and dangerous (for example, metal detectors miss plastic mines), indiscriminate and slow (for example, dogs are prone to fatigue), or expensive and hence not available to those who need them. All sensors look for some distinguishing feature that differentiates the mines from their natural surrounding either by emissions, geometrical form, internal voids, size, depth, distance from each other, material, conductivity or thermal signature.

A list of detection technologies was compiled by H. Ehlers and H. G. Kruessen of Stiftung Menschen gegen Minen (MgM, The Humanitarian Foundation of People against Land Mines), a Germany-based non-government humanitarian relief organization (web page: <http://www.mgm.com>) with help from numerous organizations including the U. S. Army. They divide the sensor systems into two general types: substance-analyzing sensors and imaging sensors.

The substance-analyzing sensors are as follows: passive/active metal detection, chemical detectors, bio-sensors (animals or animal tissue reacting to explosive vapors), and nuclear detection (by various backscatter principles such as x-ray analysis, thermal neutron analysis, fast neutron activation and quadrupole resonance). Some of the most common ones that have been tested are shown in **Table 1** in roughly increasing order of cost for the sensor. There are many reference materials for these sensors and none is a clear choice for all situations; that is, a detector that works well in the desert may not be good for snowy forested mountains.

Because of their relatively lower cost, only the first two sensor systems are in common use. False positives are due to clutter in the environment causing the sensor to alarm, so one might imagine a suite of different types of sensors for each possible scenario that would allow for cross-checking. This approach clearly drives the cost up for complete systems, making it even less accessible to the people who need it. What **Table 1** does not show is the reliability of each system measured in terms of probability of detection (PD) and false alarm for every possible situation.

To assess their performance, receiver operating characteristic (ROC) curves are measured for each detector. Schematic plots of ROC curves for AP mines are shown in **Fig. 1** for metal detectors and ground-penetrating radars (GPRs). PD is plotted relative to the number of false alarms per square meter. The clear goal is to reach nearly 100% detection with no false alarms (marked as “Desired” on the plot).

Current systems are nearing this number for anti-tank mines, but published results show that AP mine detection is far from that goal. The U. S. Army counter-mines and the U. N. humanitarian mine detection goals for PD are noted on the graph (90% and 99.6%, respectively).

The LANDMARC goal, as noted on the graph, is to approach the Army and U. N. goals in incremental steps. The clear objective is to continue to move the curves farther toward the desired sensor goal. Measures of ROC values for various terrain,

environment, and mine conditions will help to track progress and determine the success of an MIR mine detector. Success of this project will be defined as having made fair and accurate measures of the ROCs for MIR, to determine if radar-based mine detection is feasible for military or humanitarian mine-clearance goals.

Imaging sensors are used with computer equipment to enhance the user's vision to reduce the clutter problem. They can be designed for surface mines (passive video, active/passive millimeter-wave

Table 1. Tested mine detection technologies and their capabilities.

Type	Detection method	Advantages	Disadvantages
Probes or sticks	Carefully probe into the ground and search for hard materials to remove	Inexpensive	Very dangerous and slow
Metal detectors	Induction coil measures weak fields caused by metal content	Sensitive to metal mines & firing pin	Clutter, cannot find non-metallic mines
Dogs	Measure trace amounts of HE in air	Very sensitive to HE and trip-wires	Subject to fatigue and masking
Ground-penetrating radar	Microwave reflectance from dielectric interfaces	50-70% PD for non-metallic AP	Clutter, resolution, expensive
Infrared	Finds thermal differences between soil and mine (or disturbed soil)	Stand-off to 10 m	Contrast difference low, expensive
Thermal neutron activation	Analyze energy distribution of neutron backscatter using 10-MeV gammas	Very sensitive to trace HE	Large, expensive, slow, hazardous




Figure 1. Receiver operating characteristic curves. These curves are indicative of trends only and will differ greatly for varying conditions such as soil type, moisture, mine type, or weather.

radar), to look at temperature differences in disturbed soil (active/passive infrared (IR), airborne LIDAR), and at conductivity differences (GPR). GPR has been attempted by many groups for over 20 years, and MgM identifies three types, as follows:

1. Narrowband GPR (<1 GHz) — suitable for deep objects (up to 100 m); requires large expensive equipment; has depth resolution worse than 1 m.
2. Wideband GPR (FM-CW up to 3 GHz) — resolution up to 10 cm at depths to 2 m (the wide frequency band produces more information); needs excessive computing capacity and small distance to ground; expensive and not mass produced.
3. Interferometric GPR — uses synthetic aperture radar (SAR) principles with the above systems to produce better 3-D resolution; needs expensive and large radar and processing hardware.

Very little work has been done in the higher frequencies (up to 8 or 10 GHz) or with extremely wide bandwidths (many gigahertz). Aside from the expense and size, GPR has always suffered from high false-alarm rates. So, an integrated radar mine-detection system stands the most chance of widespread use if the following problems could be solved:

1. clutter, causing false positives (could try higher 3-D resolution, wider bandwidth);
2. size and weight of the system (currently too big for one person);
3. cost of the sensor; and
4. computer and software complexity.

The LANDMARC mine detection project is directly addressing all of these problems with GPR and is laying the groundwork for a detection system that has the potential to make significant contributions to humanitarian demining.

MIR Mine Detection System

LLNL has technologies and expertise in electromagnetic modeling,^{10,11} radar systems,^{2,10,12,13} and signal processing¹⁴⁻¹⁸ that can be applied to assess the limits of radar mine detection and improved techniques for clutter reduction. The stated goal is for detection of small AP mines near the surface. LLNL technologies to apply to this problem are the MIR technologies, electromagnetic modeling and simulation expertise, and image processing methodologies.

MIR has optimal radar properties for AP landmine detection, but there is still development to be done to prove its reliability and speed characteristics. Since the sensors are orders of magnitude

smaller, lighter, and less expensive than current GPRs, they can be packed together so that, for the first time, radar images of what is below the ground can be made with a portable device. Making images of the routinely sub-surface at optimal radar frequencies reduces the effect of clutter because the operator can now discriminate a mine from a rock based on its shape. The potential for a first-ever fully-3-D, real-time imaging system capable of revealing imagery of volumes beneath the earth's surface would be a dramatic contribution by LLNL. Therefore, the MIR technology can directly improve the two key performance parameters of mine-clearance efforts: reliability and speed of clearance.

The approach taken in this project has been directed toward the following five areas:

1. customizing the MIR technology for the landmine detection problem;
2. developing models of the radar and clutter signatures;
3. incorporating those models into algorithms for optimal processing;
4. performing experiments to validate the theoretical models; and
5. identifying sponsors whose needs are for a radar mine detection system.

Research results are encouraging and have been attracting the attention of mine clearance experts. The goal—which we are convinced is highly likely to be achieved—is to transition MIR technology into fielded systems that save human lives and mitigate the tremendous human suffering (thousands of lost lives and/or limbs) caused by AP land mines.

MIR Technology Overview

MIR technology is an entirely new sensor concept that is based on high-speed pulsed electronics and was developed in LLNL's Inertial Confinement Fusion program.¹⁹⁻²³ The radar exhibits a combination of interesting properties, including wideband operation, very low noise-floor, extremely low power consumption, small size and low cost, array configurability, and noise-encoded pulse generation. An MIR motion sensor circuit is shown in **Fig. 2**.

Of the MIR technology family, the swept range-gate rangefinder with transmit/receive cavity-backed monopole antennas is the sensor used for mine detection. The rangefinder, or variations of it, is used for signal and image processing applications because the return signal is similar to traditional "A-scan" from other radars. From these data, advanced signal processing and detection methods extract information about the soil under inspection.

Variations on the basic rangefinder can change frequency (1 to 4 GHz, nominally), maximum range (0.3 to 10 m possible), and sensitivity. Antenna systems control a number of these parameters and have standard modular connections on more recent




Figure 2. MIR motion sensor—a complete radar system (without antennae or 9V-battery).

MIR implementations. Plots of the radar pulse (reflected from a metal plate) and its spectrum are shown in **Fig. 3**. Notice that the radar response is relatively flat in the passband of 1 to 4 GHz, and also that system noise and jitter are low.

Progress

Significant progress has been made in all five areas. It has been clear from the beginning that the clutter problem is of major importance to mine detection. That is, traditional GPR signals are susceptible to false signatures that appear to be mines or, worse, can cause the operator to miss a mine. The design of radar systems and processing algorithms in LANDMARC has been directed toward reducing clutter.

One direction to reduce clutter has been to characterize and optimize the radar signatures for mines. The current radar specifications are given in **Table 2** for a four-element MIR array. Important goals are to have very wide band signatures with low ringing. The MIR signal, reflected from a metal plate, and its spectrum are shown in **Fig. 3**. From this image, the MIR is shown to be stable and very wideband.

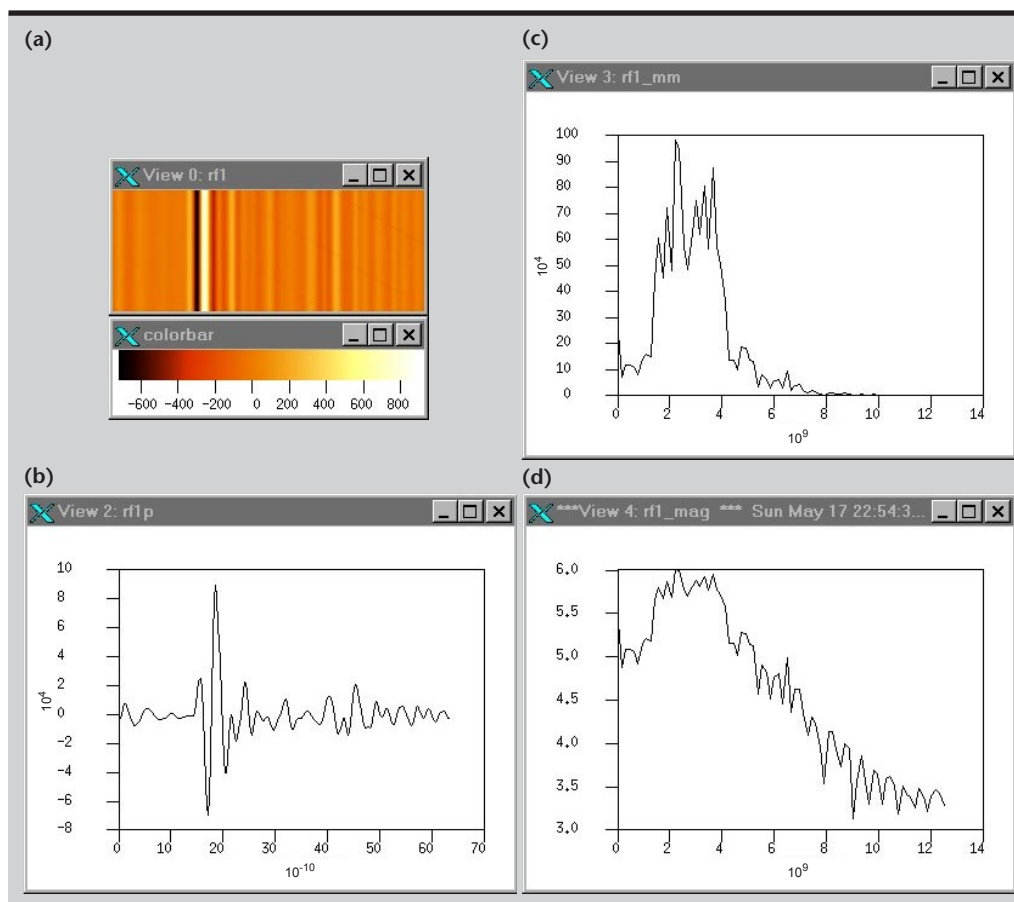


Figure 3. MIR rangefinder pulse echoed from a metal plate at 30 cm. (a) 100 waveforms stacked in an image; (b) the average of all 100 waveforms, (c) and (d) linear and log spectra of (b), respectively. The abscissa values are in seconds and Hertz; the ordinate values are in arbitrary units.

Table 2. MIR mine detection system specifications.

Parameter	Value
Radar Type:	Four-element MIR monostatic rangefinders
Frequency range	1.3 to 4.2 GHz (3dB bandwidth)
Avg. transmitted power	0.1 mW
Peak transmitted power	0.2 W
Pulse repetition frequency	5 MHz
Pulse rise time	80 ps (10 to 90% of peak)
Receiver noise figure	3 to 5 dB
Pulse width	1 ns FWHM
Timing jitter	50 ps p-p
Interface Electronics:	
Timing	Rangefinder ramp, controlled by computer
Ramp frequency	160 Hz (40 Hz for all four)
Number of averages	120
Number of range bins	256
Antenna Type:	Wideband ridged horn
Frequency range	1 to 18 GHz
Gain	3 to 10 dBi , highly-dependent on ground coupling geometry (6 nominal)
VSWR	1:1.8
Computer System:	Dual-processor Pentium, 300 MHz

Besides improving the radar itself, the LAND-MARC project has made an initial design of the signal processing performed by the system. A critical issue in this design is a clear understanding of all signal, noise, and clutter sources in the data. We define noise as any contributor to the radar measurement that degrades our capability to detect the targets of interest. With this general definition there are many contributors to the system noise model, from simple to more complex:

1. thermal receiver noise – generally modeled as white, Gaussian additive noise in the radar receiver.
2. inhomogeneity in the soil – any variations in EM properties of the soil, including density variations, moisture changes, buried rocks, or organic matter, will cause radar reflections. These will be modeled as additive non-Gaussian noise with correlations on all lengths.
3. multiplicative noise – antenna sidelobes and quantization will contribute to noise which is signal-dependent.
4. surface artifacts – because the surface is generally the strongest reflector in the returned radar signal, any imperfections in removing it can cause artifacts in the residual measurements.

5. multiple scattering – in an environment with multiple strong scatterers such as mines, rocks, or the surface the radar measurement will contain not only their primary reflections but also multiple reflections caused by the multiple possible paths back to the receiver. These signals will appear as target-like images.

We model the noise in the measurement statistically. The noise model can be most generally described by an N-dimensional joint probability density, where N is the total number of samples in the measurement. We consider only two aspects of this general distribution: the marginal point density function and the correlation of the signal in time (and later in 3-D space).

For the Gaussian case, this provides a complete description; for the non-Gaussian case, we will allow the point distribution to vary but still consider only second-order correlations. This limitation may be relaxed later if analysis of test data shows that higher-order correlations are important.

The development of noise models involves two main activities: the first is the development of a database of measurements from which statistical models can be developed. The noise database must cover the system and environmental parameter space described above for signal models. The second activity is the analysis of the noise data to extract models of point statistics and correlations

followed by the development of noise simulations based on these statistical models.

The signal processing chain, shown in **Fig. 4**, shows the planned steps for reducing clutter based on models of the noise components. The general approach to detection is based on the development of both target and noise models. These models are used to compute a likelihood that the measurement was drawn either from the signal-present (H_1) distribution or from the signal-absent (H_0) distribution. The models will generally be adapted to the measured data and generated in real time.

In the 3-D imaging approach to mine detection, a single-aperture radar can be scanned to form a synthetic 2-D array. An image in x-y-z is formed computationally using the real array in the x direction and by a synthetic aperture in the y direction. This 3-D image is then analyzed to detect the presence or absence of mines.

The general form of the detection model for the 3-D imaging radar is the same as that of the 1-D single-aperture radar. The main difference is the addition of the image formation steps, which gain resolution in x-y and also bring in new noise sources. The detection approaches can be expanded to incorporate x-y information into the template catalogs to supplement the time-frequency information in the 1-D approach.

The image formation step is likened to a refocusing of the data, and thus demonstrates significant performance improvements with this additional information. Our principal approach to clutter rejection requires the use of 3-D radar images to distinguish the shapes of mines from natural subsurface inhomogeneities.

What we have found is that if the target object (the mine) is smaller than one quarter of the highest wavelength, reconstructive imaging provides little improvement in detection. This was verified in field tests as well as simulations. The basic signal model discussed above must include the 3-D point-spread function (PSF) of the imaging formation technique. This PSF will limit the resolution of the imaging system and can be a very complicated function of the radar system, the imaging geometry, and the target itself. For our initial analyses, we have assumed the PSF is position-invariant and consists of a specific 3-D form convolved with the target and clutter distribution.

More complex effects are being analyzed using finite-difference time-domain (FDTD) models and incorporated into the statistical noise models. An EM model of the MIR signal and its response from the ground and various objects (mines) has been developed. The model has been verified by experimental means and used to characterize clutter

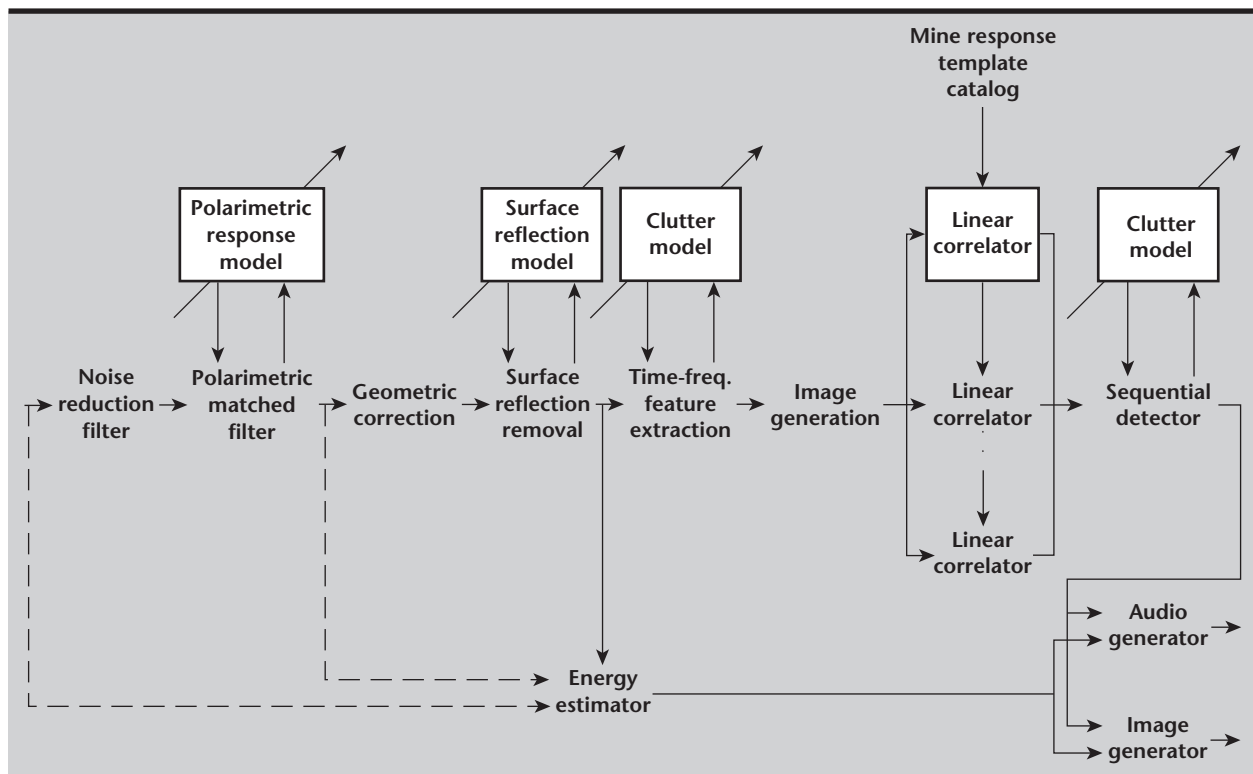


Figure 4. Signal processing block diagram for the LANDMARC mine detection system.

sources in actual data. One of the key results of the modeling effort is that the radar bandwidth for mine detection of near-surface AP mines is 0.5 to 10 GHz. This spectrum is higher than the current MIR (1 to 4 GHz).

On-going development efforts are working toward that end. Further refinements in the radar antennae and electronics have produced mine detection capabilities that are beyond the known state of radar technology. Algorithm refinements, particularly for removal of the first-surface clutter source, have provided step-by-step improvement in the signal-to-clutter ratio.

Another aspect of 3-D radar imaging is the sensitivity of the imaging to motion errors in the radar position estimates. The analysis of radar imaging parameters showed that precise positioning of radar transceivers is necessary to produce clear images that aid in clutter removal. Our initial design of a system that meets this requirement, yet is portable enough to use in the field, is a one-meter-squared array design shown in **Fig. 5**. The sensor array head is cantilevered over the area to be scanned and it produces overhead 3-D images of the sub-surface so that the operator can “see” not only the location of the buried object, but also its depth and shape. After identifying the mines in the image, the system positions its sensors over the mines for marking or removal.

Recent experiments with MIR show promise for improved AP mine detection. The developmental prototype of the mine detection system, with software for noise and clutter reduction, was assembled. A series of tests were performed at LLNL that were intended to challenge the MIR detection capabilities and provide realistic statistics for PD and false alarms.

The experiments involved unknown locations of mines (blind tests), using American-made M14 mines and intentional clutter in three soil types and in generally wet weather. Combined results of 18 independent tests of MIR compared to the Schiebel PSS-12 metal detector (commonly used at this time in most areas of the world) are shown in **Table 3**. From these results we believe that our approach is justified and should be continued.

Using laboratory prototypes and software, MIR has proven itself capable of detecting many plastic AP mine types buried in various soils. From this basic capability we plan to develop full mine detection systems that are based on the MIR concept; that is, radar-based mine detectors that are small, lightweight, low-power (battery-operated), inexpensive to replicate, and reliable.

This means that the laboratory prototypes we now have must be packaged, ruggedized, and made invariant to temperature or environmental effects. Software for control, processing, and display must be optimized and ported to a small (embedded) platform. With industry experts, other sensors must be incorporated into the system to further improve detection probability while minimizing the false alarm rate.

Extensive field testing must be done for various mines, soil conditions, and weather. Once the R&D is completed, we estimate that replication costs of such systems could be below \$10,000 when integrated by an industry partner.

Summary

The LANDMARC goals for the second year have been completed and preparations are underway for continuation into FY-99. An array of MIR transceivers has been built, mounted on a portable platform, and shown to detect buried plastic AP land

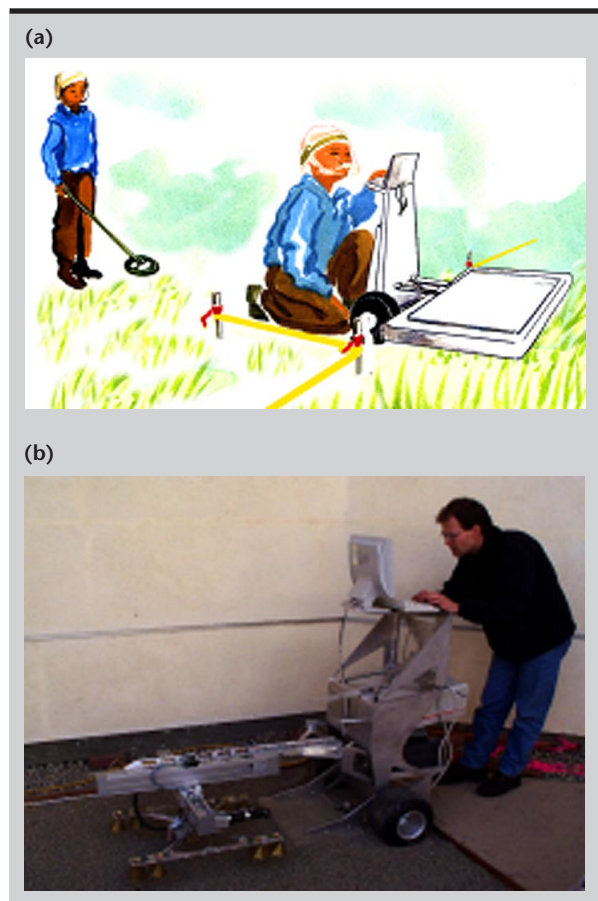


Figure 5. (a) Operational scenario for a proposed MIR mine detection system. (b) Brass-board prototype system for developmental testing.

Table 3. Results of blind experiments of mine detection using a metal detector vs results from MIR.

Average of 18 tests	Detection rate	False alarms (per detected mine)
AN/PSS-12 (metal detector)	44%	7.95
MIR (radar detector)	85%	10.02

mines using imaging software and off-line processing. Results showed this system to out-perform the metal detectors in common use today. The next steps are to further refine the system for field operations and to move toward deployment with industrial participation.

A summary of what MIR currently offers is given below.

1. MIR is the smallest, lightest, lowest power radar system known, and hence cheapest to replicate by factors of perhaps 100; however, its performance in the field under many conditions still needs to be proven.
2. Because it emits a wider frequency band than previous GPR systems, the MIR sensor alone will generally provide greater information. Depth of penetration is enough for most mine detection applications.
3. Unique 3-D imaging software developed for MIR produces underground views not currently used in mine detection. It requires accurate position information and the latest computer technology; both are achievable and could be developed. The imaging provides either horizontal slice planes or a 3-D perspective view as though peering into the ground. The array concept initiated out in this project could provide a rapid acquisition mode for this type of imaging.
4. The miniature properties of MIR make it cost-effective for arrays. MIR imaging arrays for vehicles, robots, and even hand-held systems, should improve discrimination and speed, but this needs to be verified and we are only now starting full field testing.

These characteristics are very appealing to those interested in mine detection. However, there is still much work to be done. For example, Army personnel, and others who evaluate mine detection technologies for deployment, require controlled measures of PD, false alarm rate, and scan rate in technology demonstrations. Evaluation of mine detection systems is performed by the Army in fields of their own design. These demonstrations are intended to be part of Phase 3 experiments to evaluate the system. Modeling, radar and software

refinement, and transition to deployment are also planned for the next phase.

Future Work

Recall that the overall objective of the entire project is to transition the successful MIR technology into a prototype land-mine detection system that helps mitigate the tremendous human suffering caused by small plastic AP mines. Attainment of this goal and wide dissemination of this technology will help stabilize local economies and environment, and indirectly enhance global security. Military demining objectives will also be addressed by the same system. Successful completion of this project will demonstrate improved mine detection over current technologies, enhance LLNL technical capabilities in radar imaging for other applications, and also attract cooperative participation from the U. S. Army and other sources.

Specific steps to be taken in the next phase are as follows:

1. Incorporate a higher-frequency radar into the mine detection prototype. The goal is to develop a radar with an expected 9-GHz passband (at 3 dB). A wider bandwidth will significantly improve the radar imaging capability.
2. Refine the noise and clutter models to reflect the wider passband and predict performance in various conditions.
3. Apply the improved models to the algorithms for processing, reconstruction, and detection. In particular, use first-surface echo removal to reduce clutter and study super-resolution methods through the University of California at Davis.
4. Perform the series of experiments necessary to validate the above models and measure the detection performance improvement.
5. Pursue a path toward commercialization.

We also plan to demonstrate the system for various organizations that have already expressed interest in these sensors, including industry, government agencies (DOD, State Department, DARPA), foreign governments and institutions (Japan, Norway), and humanitarian relief organizations (World Bank, U. N.).

The technologies developed for this project will have utility well beyond the life of this project. LLNL and other federal agencies have a need for radar imaging systems that have similar characteristics. Bridge inspection for the Federal Highways Administration (with whom a 64-element array has already been developed) will benefit from the mine detection technologies.


Other imaging for nondestructive evaluation of concrete, wood, and other materials is also underway. The experience and technology gained through the MIR mine detection project will significantly enhance the understanding of radar imaging and will benefit many new areas.

Acknowledgments

The authors acknowledge the contributions of the scientists, engineers and technicians involved with the MIR Project. These include J. Brase, M. Bujak, R. Cavitt, G. Dallum, R. Gilliam, G. Governo, H. Jones, M-L. Liu, J. Mast, S. Nelson, T. Rosenbury, R. Stever, M. Vigers, and P. Welsh. Administrative help from B. Wooton, M. McInnis, and R. Sachau is appreciated. We also acknowledge the continued support from the Laser Programs Directorate (J. Brase, S. Vaidya, and E. M. Campbell) and the Electronics Engineering Department (S. Dimolitsas and R. Twogood). The AN/PSS-12 metal detector was lent to us by P. Harben of LLNL.

References

1. Anderson, C., W. Aimonetti, M. Barth, M. Buhl, N. Bull, M. Carter, G. Clark, D. Fields, S. Fulkerson, R. Kane, C. Lee, F. Lee, B. McKinley, J. Page, F. Roeske, Jr., T. Rossow, P. Sargis, J. Scarafioti, P. Schaich, S. Sengupta, and R. Sherwood, (1994), "LLNL Electro-optical Mine Detection Program," Lawrence Livermore National Laboratory Internal Report, Livermore, California (UCRL-ID-118672).
2. Azevedo, S. G., D. T. Gavel, J. E. Mast, and J. P. Warhus, (1995), "Statement of Capabilities: Micropower Impulse Radar (MIR) Technology Applied to Mine Detection and Imaging," Lawrence Livermore National Laboratory, Livermore, California (UCRL-ID-5366).
3. Azevedo, S. G., D. T. Gavel, J. E. Mast, E. T. Rosenbury, and J. P. Warhus (1996), "Arrays of Micropower Impulse Radar (MIR) sensors for Subsurface Detection," *Proceedings of the EUREL Conference on the Detection of Abandoned Land Mines, IEE Conf. Pub. No. 431*, Edinburgh, Scotland, UK.
4. Azevedo, S. G., D. T. Gavel, J. E. Mast, and J. P. Warhus (1995), "Land Mine Detection and Imaging Using Micropower Impulse Radar (MIR)," *Proceedings of the Workshop on Anti-personnel Mine Detection and Removal*, Lausanne, Switzerland, pp. 48-51.
5. Clark, G. A., J. E. Hernandez, S. K. Sengupta, R. J. Sherwood, P. C. Schaich, M. R. Buhl, R. J. Kane, M. J. Barth, and N. K. DelGrande (1992), "Computer Vision and Sensor Fusion for Detecting Buried Objects," *IEEE Proceedings of the 26th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, California, p. 466.
6. Clark, G. A., S. K. Sengupta, W. D. Aimonetti, F. Roeske, J. G. Donetti, D. J. Fields, R. J. Sherwood, and P. C. Schaich (1995), "Computer Vision and Sensor Fusion for Detecting Buried Objects," *Proceedings of the Symposium of Autonomous Systems in Mine Countermeasures*, Monterey, California, April 4-7.
7. Eimerl, D. (1997), "Land-Mine Policy in the Near-Term: A Framework for Technology Analysis and Action," Center for Global Security Research, Lawrence Livermore National Laboratory, Livermore, California (UCRL-ID-130849, CGSR-98-002), August.
8. Gavel, D. T., J. E. Mast, J. Warhus, and S. G. Azevedo (1995), "An Impulse Radar Array for Detecting Land Mines," *Proceedings of the Autonomous Vehicles in Mine Countermeasures Symposium*, Monterey, California, pp. 6-112.
9. Sargis, P.D., F.D. Lee, E. S. Fulkerson, B.J. McKinley, and W.D. Aimonetti (1994), "Ground-Penetrating Radar for Buried Mine Detection," *SPIE Vol. 2217, Aerial Surveillance Sensing, including Obscured and Underground Object Detection*, April, pp. 4-6.
10. Warrick, A. L., S. G. Azevedo, and J. E. Mast, (1998), "Prediction of Buried Mine-like Target Radar Signatures using Wideband Electromagnetic Modeling," *Proceedings of SPIE Conference on Detection and Remediation Technologies for Mines and Mine-Like Targets III*, 3392, pp. 776-783.
11. Nelson, S. D., (1994), "Electromagnetic Modeling for Target-Rich Embedded Environments," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California (UCRL-53868-93), March.
12. Azevedo, S. G., J. E. Mast, S. D. Nelson, E. T. Rosenbury, H. E. Jones, T. E. McEwan, D. J. Mullenhoff, R. E. Hugenberger, R. D. Stever, J. P. Warhus, and M. G. Wieting (1996), "HERMES: A high-speed radar imaging system for inspection of bridge decks," *Nondestructive Evaluation Techniques for Aging Infrastructure and Manufacturing, SPIE Vol. 2946*, pp. 195-204.

13. Warhus, J. P., J. E. Mast, E. M. Johansson, and S. D. Nelson (1994), "Advanced Ground-Penetrating Radar," *Advanced Microwave and Millimeter Wave Detectors*, SPIE Vol. **2275**, pp. 22-24.
14. Mast, J. E. (1993), *Microwave Pulse-Echo Radar Imaging for the Nondestructive Evaluation of Civil Structures*, Ph.D. thesis, University of Illinois at Urbana-Champaign.
15. Mast, J. (1998), "SIGOP Manual," Lawrence Livermore National Laboratory Internal Report, (to be published).
16. Mast, J. (1998), "GPR Manual," Lawrence Livermore National Laboratory Internal Report, (to be published).
17. Mast, J. E., and S. G. Azevedo (1995), "Applications of Micropower Impulse Radar to Nondestructive Evaluation," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California (UCRL-53868-95), March.
18. Mast, J. E., and E. M. Johansson (1994), "Three-Dimensional Ground-Penetrating Radar Imaging Using Multi-Frequency Diffraction Tomography," *Advanced Microwave and Millimeter Wave Detectors*, SPIE Vol. **2275**, pp. 25-26.
19. Azevedo, S. G., T. E. McEwan, and J. P. Warhus (1996), "Microradar Development," *Engineering Research, Development and Technology: Thrust Area Report*, Lawrence Livermore National Laboratory, Livermore, California (UCRL-53868-95), pp. 6-17.
20. Azevedo, S. G., and T. E. McEwan (1996), "Modular MIR," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California (UCRL-53868-96).
21. Azevedo, S. G., and T. E. McEwan (1996), "Micropower Impulse Radar," *Science and Technology Review*, Lawrence Livermore National Laboratory, Livermore, California (UCRL-52000-96-1/2), pp. 16-29.
22. "The Pocket Radar Revolution," *New Scientist* (August 1995).
23. "Radar on a Chip: 101 Uses in your Life," *Popular Science* (March 1995). 



MAN-3D: A Software Tool-Kit for 3-D Image Analysis

Sailes K. Sengupta

Electronics Engineering Technologies Division

Electronics Engineering

In this project a software tool-kit is written to facilitate 3-D image data analysis. It consists of subroutines written in the commercially available image processing language IDL. The emphasis is on three aspects of image analysis: detection, identification and quantification of regions of interest in image data. All three are important topics in the analysis of 3-D data arising in nondestructive evaluation, including data from x-ray computed tomography and from ultrasonic, infrared, radar, NMR and other scans.

Introduction

Interpretive signal and image analysis plays a key role in nondestructive evaluation (NDE). Especially important are the analyses of 3-D images resulting from the use of such diverse probes as x rays, ultrasound and NMR. Typically, the data consist of a stack of 2-D image slices with each slice representing a certain physical measurement of a horizontal section of the 3-D object under investigation. Researchers are interested in locating, as well as identifying any "anomalies" or "defects" present in the specimen. We are also interested in quantifying (finding the number and size, for example) these anomalies or defects.

As a result, we are led to various questions related to the 3-D images of the objects under study. Some of these questions are negotiable in a straightforward manner akin to those involving 2-D imagery. In a 3-D context however, there are questions that are not directly answered by the methods used for the 2-D images. In some cases, the extensions, when possible, are computationally expensive and must be appropriately modified to be practically applicable.

This is true, for example, for the notion of connectedness that plays a significant role in image segmentation. In the 2-D case, we have 4-, 6- and 8-connectedness. These notions are well defined.^{1,2}

But how does one define connectedness in a 3-D image? If k is the number of voxels adjacent to a given voxel, a k -connectedness can be defined for k equaling a number ranging from 6 to 26, and subject perhaps to some 3-D symmetry consideration. Choosing k consistent with a given application is often a nontrivial task.

For another example, second-order statistics (joint distribution of 2 spatially separated pixels) plays a crucial role in texture discrimination of a 2-D image. If we try to extend it in a natural way to the 3-D case we are led to second-, or perhaps "third-order statistics" involving, respectively, the joint distribution of 2 or 3 spatially separated voxels.

Computational complexity, however, precludes the blind use of third-order statistics even when we restrict our consideration to a spatial separation of 1 voxel from one another. Even in the second-order case, one has to be selective in limiting the choice of the direction and magnitude in the choice of the voxel separation.

The examples above point to the need for a 3-D data analysis paradigm that relies on algorithms that are computationally "affordable" and application-specific in scope. The software tool-kit described in this work is based on such a paradigm. As a result, image-processing algorithms that are "global" have been preferred over the "local" ones whenever feasible, often resulting in a faster turn around time.

The software is written in the language IDL³, which is a C-like language providing access to subroutines that allow many fast and useful image-processing tasks.

In the next section we describe progress in the development of the tool-kit with reference to the specific problems in 3-D image analysis that have been addressed in this work. We also discuss a few applications from projects of recent interest at Lawrence Livermore National Laboratory.

Progress

A NDE procedure frequently results in 3-D digitized images to be analyzed and interpreted. Among the basic tasks related to such analyses are detection, identification and quantification of anomalies or defects, or more generally “regions of interest” (ROI) in the item probed. The detection of a ROI typically involves the detection of a pattern of pixels/voxels with characteristics that may be either known or unknown.

For example, the detection of regions of a given shape (rectangles, circles, or spheroids in 3-D), or known dynamic range of gray levels, or known textural characteristics in an image, is a detection of the first kind. On the other hand, the detection of all anomalous regions in an image where the pixels/voxels differ significantly from those in the background (regional or global) is a detection of the second kind.

In the subsection on detection we have concentrated on detection of unknown or partially known ROIs, relegating the treatment of the first kind to the subsection on identification. We will see there that this task is handled by some traditional methods of

classification/pattern recognition where classifiers are trained to learn the image characteristics of the ROIs based on training samples of known categories. Later, an unknown sample is presented to the trained classifier for classification.

Once the ROIs are detected and identified, we need to study them quantitatively.

For example, say we have localized and identified the types of grains of certain known textural characteristics in an image. The questions include: What is the size distribution of this particular grain type in the image? Is this correlated with the size distribution of another grain type in the image? What is the extent of this correlation, if any? Questions of this type have been addressed in the subsection on quantification.

A summary of the main subroutines developed in this project is given in **Table 1**.

Detection

As indicated above, detection usually involves localization of the ROIs when they exist. There are various means of achieving this. We have dealt with some of these methods, described briefly in the following paragraphs.

Table 1. Partial listing of main subroutines in the tool-kit for 3-D image analysis.

Detection:		
Thresholding:	histo_analyze.pro threshold.pro	histo_seg.pro
Clustering and Neighborhood Statistics:	clust_wts.pro* nabor2.pro nabor3.pro nabor3new.pro nabor_stat.pro	cluster.pro* nabor2d.pro nabor3d.pro nabor_mean_range.pro
CFAR Algorithm:	cfar_1d.pro	cfar3d.pro
Segmentation by Region Merging:	reg_merge.pro	reg_merge3d.pro
Connected Component Analysis:	conn_comp.pro	conn_comp3d.pro
Segmentation by Watershed (Meyer):	ws_meyer.pro	ws_marker.pro
Identification:		
Neighborhood Statistics:	See <i>Detection</i>	
Shape and Size Features:	See <i>Granulometry</i>	
Feature Subset:	See directory “Classify”	
Classification Routine:	See directory “Classify”	
Granulometry:		
Linear Granulometry:	lin_gran.pro	granulo_linear.pro
Granulometry by Opening:	granulo_open.pro granulo_open_line.pro	granulo_open_general.pro

* Indicates IDL subroutines

Global Thresholding. When the intensity level in the ROI is known to have a threshold value (upper or lower), or known to fall in an interval, global thresholding is a powerful way of localizing it. The threshold limits and the interval boundaries can sometimes be provided by the physical properties related to the gray-level/intensity.

For example, in x-ray radiography, tomography, ultrasound, and infrared imagery, the property could be intensity, x-ray attenuation coefficient, sound attenuation, or surface temperature, respectively. Very often however, the limits have to be inferred from the statistical distribution of the gray levels of pixels/voxels in the image. Specifically, these limits are determined from the “peaks” (or modes) and “valleys” (separating two consecutive modes) of the histogram of the distribution when they are prominent.

Global Clustering. Frequently, intensity of the image pixels/voxels alone does not provide the desirable separation between the ROIs, the background, and other objects. In such cases, one may want to pursue this separation in a higher-dimensional space in a multivariate setting. In addition to the intensity, one might include other “features” of the pixels/voxels in the image, such as some statistical parameters associated with a suitably defined neighborhood of them.

More generally, elements of the feature vector could represent gradient magnitude, gradient phase, intensities at multiple spectral bands, or other attributes. It may also be some other measured property of the pixel/voxel. The separation of the ROIs from the background is then sought in the multi-dimensional feature space by an appropriate, unsupervised training procedure such as “clustering.” Pixels/voxels that are “close” to one another in the feature space as measured by their Euclidean or some other distance, are considered “similar.” In a clustering procedure, these similar objects are clustered together. In this work we have used the K-means clustering algorithm.^{4,5}

Constant False Alarm Rate (CFAR) Algorithm. When the intensities in an image vary systematically, with a possible trend, the standard thresholding technique is unlikely to work. In such cases, detection can sometimes be implemented satisfactorily by means of the CFAR algorithm. This is particularly effective when the local variance of the intensity changes at a smaller rate than the local mean.

In this algorithm, one first subtracts the “local” mean from each pixel/voxel. Then a statistical test is applied to each pixel/voxel to determine if it is different “significantly” from its background pixels/voxels.

Although computationally intensive, the algorithm allows for a pre-assigned false alarm rate in the detection of ROIs.⁶

Region-Merging-Based Segmentation and 3-D Connected Component Analysis. In this method, each 2-D slice is segmented by a region-growing/merging method with a label attached to each segment by means of a connected component labeling algorithm. Adjacent connected components in consecutive slices are then put together and re-labeled to complete the 3-D connected component analysis.^{7,8}

Segmentation by Watershed and 3-D Connected Component Analysis. This is the same as above, except the 2-D segmentation is done by the watershed algorithm.⁹

Identification

As indicated in an earlier paragraph, identification is frequently achieved by a “supervised learning” procedure based on “ground truth” and is outlined in the following steps. The “ground truth” here refers to the known classification categories of the training pixels/voxels.

Feature Computation. Various features are computed for “training” samples of known object types. The nature of features varies widely from one application to another. In this work we have provided routines to compute 2- and 3-D neighborhood statistics and extraction of first- and second-order statistics from these neighborhood statistics.

Feature Subset Selection. If the number of features is large, a smaller subset of these features containing most, if not all, class discriminatory information is often desirable. Frequently, due to resource constraints, the experimenter has a small sample to work with, and as a result is not able to estimate accurately the distribution parameters in a higher-dimensional feature space.

In a subset selection procedure, feature subsets are selected based on their “separability measure,” which provides an index of how well the object types in the feature space consisting of a given subset of features are separated.¹⁰

The simplest possible separability measure in a multi-dimensional feature space is based on the Euclidean distance between pairs of points representing the selected features of samples from two populations. This applies when features may be assumed independent. The minimum of such pairwise distances after appropriate normalization gives the separability measure of the subset of features under consideration. When the features

are correlated, one can replace the Euclidean distance by other distances, such as Bhattacharya or Matusita.

Training and Classification. A suitably chosen “classifier” is then trained, based on the features of the training sample. A classifier is a system that creates class partitions in the feature space based on labeled (known) samples. Based on the features computed from a new sample, the trained classifier checks which cell of the partition the new sample belongs to and labels it accordingly. We have chosen a non-parametric Bayes classifier for the present work. It is also termed the “Probabilistic Neural Network” (PNN). The trained classifier is then used to classify a sample as needed.¹¹

Quantification

Having localized and identified the ROIs in the image, the researcher would like to quantify the ROIs in an appropriate manner. We have used the technique of morphological granulometry^{12,13} to implement this capability. Traditional granulometry is an essential tool of quantitative morphological image analysis. It consists of sequential application of openings (closings) with structuring elements that are monotonic.

The principle of granulometry can be expressed in simple terms as follows. Imagine a large collection of “grains” of different sizes, filtered through sieves of progressively increasing mesh sizes. Numerical characteristics of the residues at each stage are measured until nothing remains. The successive residues provide us with a quantitative measure of the grain size distribution.




Figure 1. Representative slice of the CT scan of an asphalt sample containing asphalt, rocks and air bubbles.

In fact, by varying the shapes and orientations of the mesh elements in the sieves, it is possible to get a complete quantitative morphological characterization of the grains. The method was first conceived by Matheron who called it “morphological granulometry.”

As a practical application, this method provides a direct computation of the size distribution of “grains”/“pores” in a binary image in one, two, or three dimensions without having to do a connected component analysis first. An interesting feature in this implementation is that one can detect the size distribution of the image grains along any specific direction, or more generally, those fitting specifically shaped objects with a given spatial orientation.

Applications

In this section we provide four typical applications of the subroutines in our tool-kit. This will include 3-D images from computed tomography using x-ray and ultrasonic imagery.

1. Asphalt sample. The Federal Highway Commission provided us with an asphalt sample of approximate diameter 102 mm. The objective was to quantify such components as tar, aggregate, and porosity in the sample. A linear-array detector CT scanner (LCAT) operating at a maximum tube potential of 160 keV (1.9 mA) was used, with the data filtered intentionally to reduce beam-hardening to acquire a 3-D CT image. The image (Fig. 1) shows a cross-section measuring 102 mm with a pixel size of 225 μm and a slice plane thickness of 500 μm . Figure 2 shows the attenuation histogram of the slice.




Figure 2. Histogram of the pixel attenuation in the CT image shown in Figure 1.

The image is thresholded using the “valley” points $v = [0.04, 0.2]$, resulting in the three “segments” shown in **Figs. 3, 4, and 5**. In **Fig. 3**, attenuation less than 0.04, represents the voids and background; attenuation between 0.04 and 0.2, represents the asphalt; and the high attenuation voxels represent the inclusions. Pixels that are near zero represent air both inside and outside the sample. Thus this simple method has separated air from the rest of the sample.

The first segment (**Fig. 3a**), shows that the background (in white) and the air pockets (also in white) in the block have been separated. The second segment (**Fig. 4b**) shows the rest, including a water droplet at the bottom but excluding the speckled bright spots evident in **Fig. 1**. The third segment (**Fig. 4c**) shows the speckles.

The histogram analysis (*histo_analyze*) and segmentation (*histo_seg*) routines used in this example have been used on the 2-D slice but they are equally applicable with similar results to the global 3-D data.

While thresholding separates some of the objects, it fails when there are physically different objects with highly overlapping ranges of intensities. This is the case with the asphalt, rocks, and water vial in this image. To address this problem we try to separate the objects in a higher-dimensional “feature” space consisting of a set of neighborhood statistics of each

pixel/voxel. Specifically, we choose a window centered at each pixel/voxel and compute a set of statistics for the pixels/voxels in the window. We then cluster the pixels/voxels (inclusions excepted) of the image into a fixed number, K , of clusters, where K is the number of separable objects known to be present in the image. The results of the clustering using $K = 3$ are shown **Fig. 4**.

A tabulation of the percentages of pixels representing the four different objects within the aggregate is given in **Table 2**.

2. Aluminum part. A second example demonstrates the use of the CFAR 3-D code for isolating regions suspected to be porous in 3-D imagery obtained using a high-frequency ultrasonic probe on an aluminum part. One of the objectives in the porosity study is to get an estimate of the pore-size distribution. **Figures 5a to d** show a representative slice in the middle part of the ultrasonic scan of the aluminum part. **Figure 5a** shows the raw image. **Figure 5b** shows the result of a pre-processing step,

Table 2. Composition of the aggregate sample by percentage.

Rock	49
Asphalt	29
Void	16
Inclusion	6




Figure 3. Binary image showing the three segments obtained by using the threshold vector $[0.04, 0.2]$ in the image in Figure 1.




Figure 4. Three segments showing the asphalt, rocks, and the water droplet.

the absolute value of the intensity deviation from the mean image intensity. **Figures 5c** and **d** show the result of applying the CFAR 3-D code to isolate the “anomalies” using a window size of 11×11 and 15×15 , respectively, with the false alarm rate set at the five percent level.

3. A third example illustrates other routines for object detection. We used a combination of the CFAR and a 3-D connected component analysis for a 3-D segmentation. The data used are taken from the Active and Passive Computed Tomography (A&PCT) program for waste drum radioactivity assay. The objective here is to localize the regions of high radioactivity. The frames in **Fig. 6** represent $14 \times 14 \times 17$ slices of A&PCT data acquired at 414 keV. Slices 1 to 15 arranged from top to bottom and left to right in three rows are shown in **Fig. 6**.

The light blobs indicate the presence of high radioactivity. We applied the CFAR_3D algorithm to localize the radioactive objects in the image. The result is shown in **Fig. 7**.

A 3-D connected component analysis led to three distinct segments (not shown) consisting of 6, 16, and 44 voxels, respectively.

4. Our final example is taken from the porosity study of a die-cast pre-stressed aluminum (A-356) tensile bar, to illustrate how methods of morphological granulometry can be used to derive the 1-, 2- and 3-D size distribution of pores (or grains) using 3-D tomographic data. These in turn can be used to model metal composite behavior under different kinds of stress.

Figure 8 shows a single slice of a 3-D CT scan of the aluminum sample using the PCAT system with 80 keV. The CT image has a pixel-resolution

Figure 5. (a) High-frequency ultrasound c-scan image of an aluminum part. Successive steps used to isolate the pores include: (b) absolute deviation from the mean, a preprocessing step; (c) results of the application of CFAR with a 3-D kernel of 11 voxels to (b); and (d) results of the application of CFAR with a 3-D kernel of 15 voxels to (b).





Figure 6. Slices 1 to 15 (top left is blank) of A&PCT data acquired at 414 keV.



of 20.5 μm . It is of interest to find the size distribution of the pores in the matrix along a chosen direction, in particular along the three axes.

One may also be interested in finding the 2- and 3-D size distribution of these pores. The granulometry routines developed in this work have been used to find such distributions applied to the thresholded binary image obtained from the original. **Figures 9** and **10** show the size distributions along the x- and z-axis, and **Fig. 11** the same for the 3-D volume. If needed, a comparison of the two size distributions can be made at this stage.

Future Work

Three-dimensional image analysis resulting from a variety of probes such as x-ray CT, A&PCT, ultrasonic scans, and NMR present a formidable challenge to the data analyst. In this work we have provided a tool-kit based on the commercially available software IDL to address some of the interesting classes of problems that arise in such analyses.

Several possible extensions to improve and extend the capabilities of this tool-kit are being planned. First, like all useful tool-kits, there is the need for a graphical user interface (GUI). The GUI




Figure 7. The localized highly radioactive regions in slices 1 to 15.




Figure 8. Representative CT slice of a pre-stressed A-356 die-cast aluminum tensile bar with pores showing in dark shade.




Figure 9. Pore-size distribution along the x-axis in the tensile bar




Figure 10. Pore-size distribution along the z-axis in the tensile bar.

allows the investigator to concentrate on the data analysis and inference process without getting bogged down in the image-processing details.

Second, some of the algorithms used for granulometry in this tool-kit are of complexity $O(n)$ or higher, where n is the number of pixels/voxels in the input image. Some recent developments in morphological image processing on “fast” openings^{14,15} will allow us to replace these slow executing algorithms by faster ones.

Third, we would extend single granulometry in this work to “multivariate granulometry.” A single granulometry yields the size distribution of a single variable where structuring elements grow at the same rate so that their relative sizes are constant. Multivariate granulometry¹⁶ allows each structuring element to grow at its own chosen rate, independent of the scales of other structuring elements enabling better representation of the textural features. A texture classification protocol using multivariate granulometry will capture the morphological characteristics of an image better and allow more accurate image classification.

Acknowledgments

I am grateful to C. Logan, P. Roberson, D. Chinn, and J. Haskins for supplying the data shown in the applications above. I am particularly grateful to H. Martz for carefully reviewing an earlier draft and suggesting numerous improvements.




Figure 11. Pore-size distribution using openings by boxes.

References

1. Rosenfeld, A., and A. Kak (1982), “Digital Picture Processing,” 2nd ed., Vols. 1 and 2, Academic Press, Orlando, Florida.
2. Serra, J. (1982), “Image Analysis and Mathematical Morphology,” Academic Press, New York, New York.
3. IDL Version 5.0 (1997), Research Systems Inc., Boulder, Colorado, March.
4. Ball, G. H., and D. J. Hall (1966), “ISODATA, A Novel Method of Data Analysis and Pattern Classification,” *International Communication Conference*, Philadelphia, Pennsylvania, June.
5. Jain, A. K. (1989), “Fundamentals of Digital Image Processing,” Prentice Hall, Englewood Cliffs, New Jersey.
6. Reed, I. S., and X. Yu (1990), “Adaptive Multi-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution,” *IEEE Trans. on Acoustics, Speech and Signal Processing*, June.
7. Levine, M. D. (1985), “Vision in Man and Machine,” McGraw-Hill, New York, New York.
8. Haralick, R., and L. Shapiro (1992), “Computer and Robot Vision,” Vol. 1, Addison-Wesley, Reading, Massachusetts.
9. Beucher, S., and F. Meyer (1993), *The Morphological Approach to Segmentation: Watershed transformation* in “Mathematical Morphology in Image Processing,” E. Dougherty, ed., Marcel Dekker Inc., New York, New York.


10. Fukunaga, K. (1990), "Introduction to Statistical Pattern Recognition," 2nd. ed., Academic Press, San Diego, California.
11. Specht, D. F. (1990), "Probabilistic Neural Networks," *Neural Network*, Vol. **3**, No. 1, pp. 109–118.
12. Serra, J. (1982), "Image Analysis and Mathematical Morphology," Academic Press, New York, New York.
13. Dougherty, E., J. Pelz, F. Sand, and A. Lent (1992), "Morphological Image Segmentation by Local Granulometric Size Distributions," *Journal of Electronic Imaging*, Vol. **1**, (1), January.
14. Crespo, J., R. W. Schafer, J. Serra, C. Gratin, and F. Meyer (1997), "The Flat Zone Approach: A General Low-level Region Merging Segmentation Method," *Signal Processing*, Vol. **62**, pp. 37–60.
15. Vincent, L. (1994), "Fast Opening Functions and Morphological Granulometries," *SPIE*, Vol. **2300**, pp. 253–267.
16. Batman, S., and E. Dougherty (1997), "Size Distributions for Multivariate Morphological Granulometries: Texture Classification and Statistical Properties," *Optical Engineering*, Vol. **36**, (5), pp. 1518–29. 



Image Recovery Techniques for X-Ray Computed Tomography in Limited-Data Environments

Dennis M. Goodman and Jessie A. Jackson

*Laser Engineering Division
Electronics Engineering*

Maurice B. Aufderheide

Defense Systems

Erik M. Johansson

*Consultant
Tracy, California*

There is an increasing requirement throughout Lawrence Livermore National Laboratory (LLNL) for nondestructive evaluation using x-ray computed tomography (CT). Conventional CT methods are non-iterative algorithms that have the advantage of low computational effort, but are not sufficiently adaptable to incorporate prior information or non-Gaussian statistics. In many cases, restrictions on data acquisition time, imaging geometry, and budgets make it infeasible to acquire projection data over enough views to achieve desired spatial resolution using conventional CT methods. Most currently existing iterative tomography algorithms are based on methods that are time-consuming because they converge very slowly, if at all. The goal of our work was to develop a set of limited-data CT reconstruction tools and demonstrate their usefulness by applying them to a variety of problems of interest to LLNL. In the second and final year of this project we continued our development of reconstruction tools and have demonstrated their effectiveness on several important problems.

Introduction

Tomography is used throughout LLNL to investigate the composition and structure of objects noninvasively. Conventional CT algorithms have the advantage that they are noniterative, hence, they are computationally efficient, but they suffer from a lack of flexibility in that they can't incorporate prior information about the solution, nor can they use other than the squared error criterion to fit the data.

Because the image recovery problem associated with most tomography applications is ill-posed, which is particularly true for limited-data situations, all possible prior information about the unknown object must be used to produce high-quality images. Furthermore, although least-squares (that is, minimizing a sum of squares error function that measures the mismatch between the data and the model) is appropriate if the data has a Gaussian distribution with known constant variance, using it when the data has a non-Gaussian distribution can seriously degrade the quality of reconstructed images.

An example is counting problems (such as emission tomography) where the data consists of particle or photon counts that typically have a Poisson distribution. Another example is outlier-corrupted data, in which the distribution function is not known exactly, but it is known that the data pixels are subject to outliers that occur infrequently, but greatly distort those pixels where they occur. These outliers result from a variety of problems including a few bad detectors in a CCD array, improper assumptions about the model, or "hits" from extraneous radiation or particles. Outliers have highly deleterious effects on the reconstructed image when squared error is used.

The problems our techniques have been designed to handle can be described as follows. First there is a linear equation that models the relationship of the unknown to the data:

$$\hat{\mathbf{y}} = \mathbf{A}\hat{\mathbf{x}}, \quad (1)$$

where the vector $\hat{\mathbf{x}}$ represents the unknown image we wish to reconstruct; and \mathbf{A} is a matrix that

represents such effects as geometry and absorption that lead to the expected data, $\hat{\mathbf{y}}$, if $\hat{\mathbf{x}}$ were the true image.

The actual observed data is \mathbf{y} , and the estimate of the true image is obtained by using an optimization algorithm to adjust $\hat{\mathbf{x}}$ so that $\hat{\mathbf{y}}$ is a best fit to the actual data according to some criterion function, say

$$f(\mathbf{y}, \hat{\mathbf{y}}), \quad (2)$$

where the choice of the criterion function, as we noted earlier, depends on the statistics of \mathbf{y} . Consequently, our reconstruction problem involves minimizing, with respect to $\hat{\mathbf{x}}$, a function of the form

$$L(\mathbf{y}, \hat{\mathbf{x}}) = f(\mathbf{y}, \mathbf{A}\hat{\mathbf{x}}). \quad (3)$$

This function is usually a negative log-likelihood function, so the image we recover is the maximum likelihood estimate of the unknown.¹ Because our problems typically are ill-posed, we add penalty parameters and impose prior constraints on \mathbf{x} . This leads to the minimization problem:

$$\mathbf{x} = \underset{\mathbf{x} \in \mathbf{S}}{\text{Argmin}} \left\{ L(\mathbf{y}, \mathbf{x}) + \eta \|\mathbf{x}\|_1 + \lambda \|\mathbf{x}\|_2 \right\}. \quad (4)$$

Here we are recovering a constrained penalized maximum likelihood estimate: the extra Euclidean norm and/or absolute value norm terms penalize large $\hat{\mathbf{x}}$, and the parameters η and λ determine the degree of penalty. It is also possible to include terms that penalize derivatives of \mathbf{x} , thereby imposing a larger penalty on its higher frequency components.

If outliers are present, we do not use the maximum likelihood approach directly. Instead, we select $f(\mathbf{y}, \hat{\mathbf{y}})$ to make the reconstructed image less sensitive to outliers. We give an example in the next section.

The subset \mathbf{S} represents upper and lower bounds on the components of the image vector $\hat{\mathbf{x}}$. These bounds provide crucial prior information about the solution, and can greatly improve the quality of the recovered image. Because the unknown image usually represents non-negative quantities such as absorption or energy, the most common constraint is that all the components of $\hat{\mathbf{x}}$ be non-negative. Other constraints on subsets of pixels can result from the known extent of the object and regions of known voids or occlusions.

The optimization problem we have described above is a difficult one for several reasons. The most obvious is its very high dimensionality: both $\hat{\mathbf{x}}$ and \mathbf{y} frequently have 10^6 or more components. Another reason is that bounds that we have imposed on the

solution require a constrained optimization technique. Many iterative nonlinear minimization algorithms require storing an inverse of the matrix of second partials of $L(\mathbf{y}, \hat{\mathbf{x}})$ with respect to the components of $\hat{\mathbf{x}}$. This is clearly not possible for our problems because a $10^6 \times 10^6$ matrix is too large to store.

The answer for unconstrained problems is either the conjugate gradient algorithm or the limited memory quasi-Newton algorithm. Both of these algorithms in effect store low-order approximations of the second-order information that is provided by the inverse matrix of second partials.

For bound-constrained problems, conventional optimization techniques²⁻⁴ usually allow only one variable per iteration to attain a bound, so for very large problems these techniques are very slow because they spend too much time finding bounds. Consequently, standard iterative tomography algorithms^{5,6} are based on methods that can attain multiple bounds in an iteration. Unfortunately, however, they are essentially steepest descent techniques that use no second-order information about $L(\mathbf{y}, \hat{\mathbf{x}})$ whatever, so they converge very slowly, if at all.

Progress

Optimization Algorithm Development

Our ability to solve difficult tomography problems is the result of two specialized optimization algorithms that we have developed. The first is an extension of the conjugate gradient algorithm that incorporates bound constraints on the variables.⁷ This constrained conjugate gradient (CCG) algorithm is unique in that it incorporates a bending linesearch that permits multiple bounds to be attained during a single iteration. In previous years, the algorithm was applied with great success to a variety of practical problems.⁸⁻¹⁵

The second algorithm is a limited-memory quasi-Newton (QN) algorithm that permits upper and lower bounds on the variables. This year we completed the implementation of our QN algorithm, based on the derivations in Reference 16.

The very large dimensions of \mathbf{A} typically make calculating the matrix-vector product of Eq. 1 the most computationally expensive part of finding a solution. Which algorithm is most appropriate depends on the structure of \mathbf{A} .

In particular, the value of $L(\mathbf{y}, \hat{\mathbf{x}})$ and a related directional derivative must be calculated at each new sub-iteration within the bending linesearch. Completely recalculating Eq. 1 each time is very

expensive, but in some cases a different approach is possible. This is because it is necessary to recalculate only that part of **Eq. 1** corresponding to components of $\hat{\mathbf{x}}$ that have either attained or left bounds during a linesearch sub-iteration. For example, if only one component of $\hat{\mathbf{x}}$ has changed its bound status, then it is necessary to compute only one scalar-vector product, the product of that component with its corresponding column of **A**—rather than the entire vector-matrix product shown in **Eq. 1**. If **A** has many columns, this saves considerable effort.

If the columns of **A** are readily available, then CCG appears to be the best choice. However, in many tomography problems, **Eq. 1** is not computed in the usual manner. Often **A** is a discretized version of the forward projection operator; therefore, it represents a set of line integrals, and it is very sparse. Consequently, **Eq. 1** is computed via a rule that calculates only the line integrals and wastes no time on the large portions of **A** that consist of zeros.

Another example is where **A** is the kernel of a shift-invariant blurring function. In this case, using a convolution algorithm based on the fast Fourier transform (FFT) is by far the most efficient method for computing **Eq. 1**.

In some of our future work on the Advanced Hydrotest Facility (AHF)¹⁷ we anticipate that **A** will be of the form **A** = **BC**, where **C** is a projection matrix and **B** is a blur kernel that represents the effects of detector and source geometries. In this case, both a sparse matrix rule and an FFT convolution algorithm are required. In all the cases discussed in this section, the programming effort required to calculate individual columns of **A** is considerable, and often not worthwhile.

The advantage of the QN method is that bending is used first on a low-order quadratic approximation to $L(\mathbf{y}, \hat{\mathbf{x}})$ to produce a direction for the linesearch. This greatly reduces the need for bending during the linesearch itself, yet it still permits the algorithm to attain multiple bounds during an iteration. In the QN algorithm we have included the options both to perform additional bending in the linesearch if the columns of **A** are available, and to perform no bending in the linesearch if they are not.

Robust Tomography Algorithms

This year we completed work on and demonstrated the usefulness of our robust tomography technique. The squared error function is

$$f(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^N r_i^2 = \mathbf{r}'\mathbf{r} , \quad (5)$$

where the residual vector $\mathbf{r} = \mathbf{y} - \hat{\mathbf{y}}$ is one indication of mismatch between the observed data and the model of the data.

If **Eq. 5** is the criterion function in **Eqs. 2** and **3**, then the reconstructed image defined in **Eq. 4** will be very sensitive to outliers. For example, suppose $N = 10,000$ in **Eq. 5**, and the magnitude of a typical component of the residual is 1.0, except for one component, r_k , whose value is 100.0. Then r_k contributes roughly as much to the squared error as do all the other components combined, and the minimization algorithm will try very hard to match the outlier, y_k , at the expense of all the other data points.

The solution is to use a robust criterion function^{18,19} that reduces the influence of larger residuals. The most common robust criterion is

$$f(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^N g(r_i) , \quad (6a)$$

where

$$g(r) = \begin{cases} r^2 & \text{if } |r| \leq c \\ 2c|r| - c^2 & \text{if } |r| > c \end{cases} . \quad (6b)$$

The function $g(r)$ transitions smoothly from a squared penalty to a linear penalty, thereby reducing the influence of large residuals. The choice of c is data-dependent.^{18,19}

We have demonstrated both that CCG is very effective in minimizing **Eq. 4** when **Eq. 6** is the criterion function, and that this approach greatly improves image quality in the face of outliers. We present a simulated example. The true unknown is shown in **Fig. 1**, and the resulting projection data with noise added is shown in **Fig 2**.

This data is in the form of a sinogram, a 2-D image created by displaying all the ray sums at one angle vs all angles (projections) obtained. The noise is a Gaussian mixture, with probability 0.99 that the noise at a data pixel is Gaussian with standard deviation σ , but with probability 0.01 it is Gaussian with standard deviation 100σ . Consequently, there is an outlier 1% of the time.

These outliers are obvious as speckles in **Fig. 2**. The reconstructed image using **Eq. 5** is shown in **Fig. 3**; the reconstructed image using **Eq. 6** is shown in **Fig. 4**. Outliers caused the streaks in **Fig. 3**. Our robust technique reduced the rms error between reconstruction and true object by a factor of two. Although robust statistical methods have been applied to a variety of problems, to our knowledge this is their first application to tomography. The effectiveness of CCG made this possible.

Figure 1. Original object used for CCG simulated example.





Figure 2. Sinogram of projections of object shown in Figure 1, with Gaussian mixture noise added.






Figure 3.
Reconstruction from
projection data in
Figure 2, using
squared error crite-
rion.




Figure 4.
Reconstruction from
projection data in
Figure 2, using
robust technique.

We have also successfully tried other robust criteria. Although we presented only a simulated example, many tomography problems are plagued by outliers, and this method has promise for a variety of real problems.

Waste Drum Assay

Using emission tomography to characterize mixed waste drums has been studied at LLNL for several years.²⁰ An active and passive computed tomography (A&PCT) technique has been developed that first uses an external radioactive source and active tomography to map the attenuation within a waste barrel. This attenuation map is used to define the matrix **A** for the passive or emission tomography problem that is of interest.

The emission tomography part of the waste drum problem gave us opportunities to demonstrate the effectiveness of our techniques and to develop them for general emission tomography problems.

At each detector position we acquire the entire gamma-ray spectrum, and two counts of gamma-ray emissions are taken. The first count is in the region of a spectral peak of the isotope of interest; the second in a region just outside this spectral peak. The purpose of the second measurement is to determine the level of background radiation and remove its effects on the first measurement.

In previously developed algorithms, the net counts due to the isotope were obtained by subtracting the second measurement from the first. Two maximum likelihood expectation-maximization (MLEM) algorithms were developed, UCSF-MLEM²¹ and APCT-MLEM,²² and applied to the corrected data to obtain a 3-D image of isotope activity. The sum of counts over all the image voxels is related to an estimate of isotope activity within the drum. Unfortunately, subtracting the two counts and then using a maximum likelihood algorithm on the net counts is not a correct application of the likelihood principle. Furthermore, this approach can violate physical reality because there is a non-zero probability that a net count will be negative.

The two MLEM algorithms lack the flexibility to implement the correct log-likelihood function. Last year we derived the correct log-likelihood function;²³ this year we developed a new algorithm, APCT-CCG, for the waste drum problem, and we studied its behavior on both real and simulated data.²⁴

The flexibility of CCG made implementing the log-likelihood function relatively simple. It was necessary only to select the appropriate criterion function that accounted for Poisson statistics and incorporated peak and background measurements at each detector location. **Figure 5** compares applying the three algorithms to real data from known Pu-239




Figure 5. Comparison of assay accuracy for three reconstruction methods.

sources. APCT-CCG was out-performed on only one of the seven sources. Although the voxel sum is the most important parameter for waste drum assay, the reconstructed image is more important for most emission tomography problems. APCT-CCG also did a much better job of reconstructing the image.

A test case was generated using simulated data. A simulated 3-D image was created with three slices. Each slice is 14×14 voxels. A point radioactive emission source of 30,350 counts was placed on the center slice at voxel location (5, 5), that is, just off the center of the slice. Using the system matrix, the image was forward projected to create three sinograms. A level background equal to the maximum signal strength was added to these sinograms. The level was carefully chosen to be consistent with and representative of empirical data.

The simulated sinograms were randomized by passing them through a Poisson random generator.

Another set of sinograms were created with the same background level. These background sinograms were also randomized. Both the gross and background sinograms were used as input for the APCT-MLEM and APCT-CCG codes. The results are shown in Fig. 6.

The APCT-MLEM image is spread out over 3×3 voxels within each slice and across all three slices. Its total assay yields 36,710 counts, that is, 121% of the actual value. The APCT-CCG results are more localized and its assay value of 9,936 counts is much closer to the original 30,350 counts, that is, within 1%.

The other important observation is that the APCT-CCG-code-calculated sinograms are more representative of the original source data than the source plus background, which is not the case of the APCT-MLEM sinogram results. We also found this to be the case for real data.²⁰

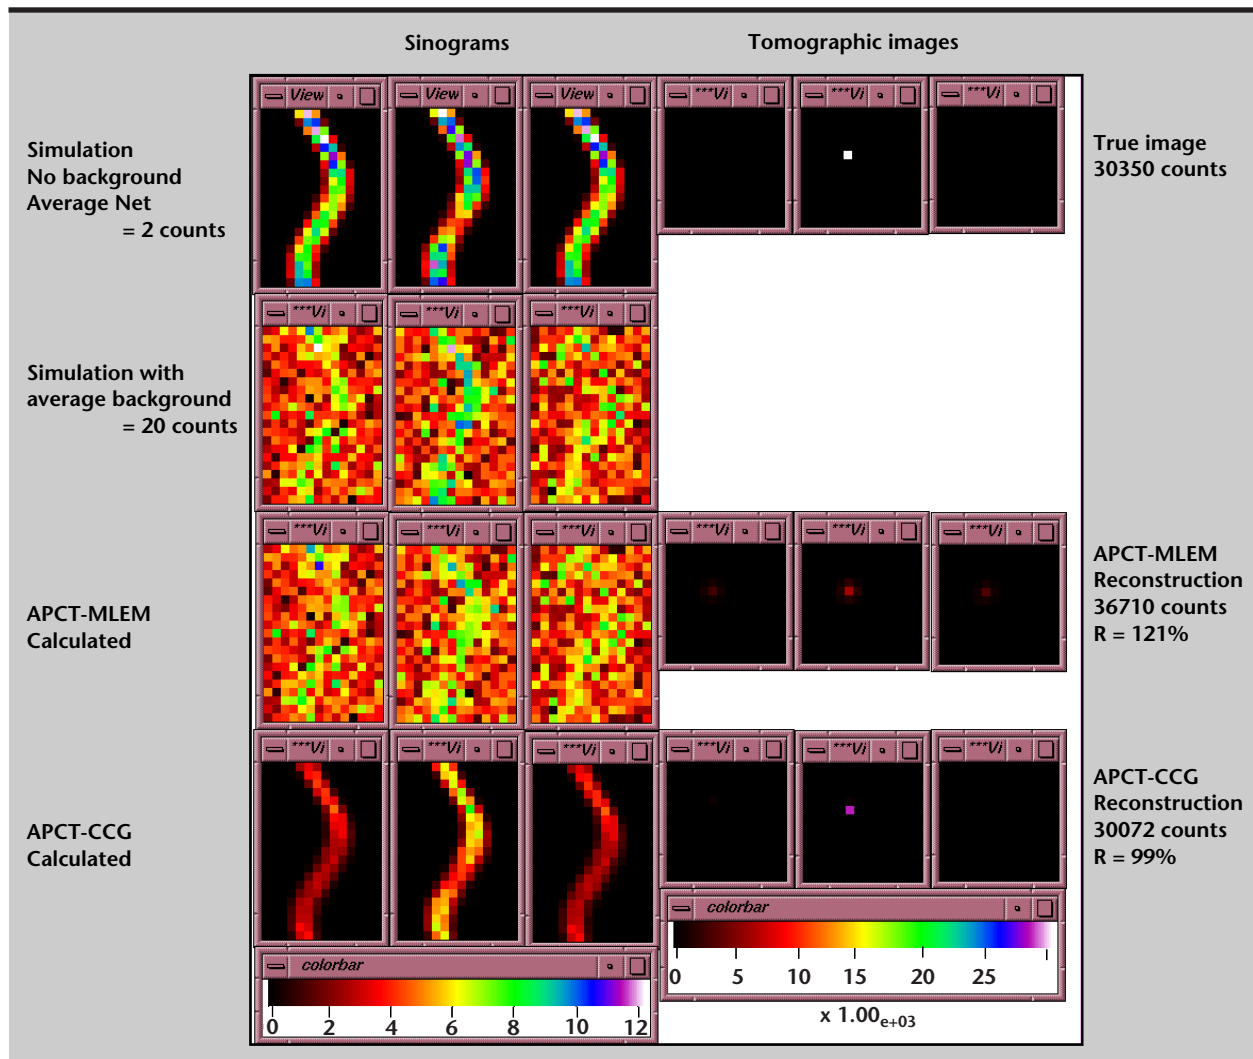


Figure 6. Simulated image and sinogram reconstruction.

For this problem it became obvious to us that model selection and validation are important issues, particularly when accurate estimates of parameters (for example, isotope activity) are as critical as good images. In the case of Gaussian noise problems where the squared error function is used, tomography can be viewed as a linear regression problem, so the usual χ^2 tests can be applied to the squared error to measure goodness of fit and to choose between models. The sinogram of the residuals is a good image to observe the adequacy of a particular model.

However, this is not the case if the noise is not Gaussian. In studying the waste drum problem, we realized that statistically it can be viewed as a generalized linear model in which there is a linear relationship between the unknown and the data, but the statistics are not necessarily Gaussian.^{19,25,26} This is basically the model we defined in the Introduction. In this context it is possible to slightly modify $L(\mathbf{y}, \hat{\mathbf{x}})$ so that it becomes a deviance function that has the same minimum with respect to $\hat{\mathbf{x}}$, but exhibits behavior that is approximately χ^2 .

Similarly, it is possible to define other sinogram images that behave as the usual residual image does for the squared error case. We have applied these ideas to the waste drum problem, and they will be useful for other tomography problems as well.

Neutron Imaging

LLNL is currently developing a high-energy (10 to 15 MeV) neutron imaging system for use as an NDE


tool in support of the Enhanced Surveillance Program (ESP). This approach to tomography promises to be a powerful technique for probing the internal structure of thicker objects that may be opaque to x rays and lower energy neutrons.

Imaging experiments using neutron radiography were conducted at the Ohio University Accelerator Laboratory (OUAL) in FY-98. The object imaged was a right-circular Pb cylinder with an outer diameter of 4 in. and a 2-in.-diameter polyethylene insert, as illustrated in **Fig. 7**. The insert was split into two half-cylinders with one serving as "blank" and the other having a series of 10-, 8-, 6-, 4-, and 2-mm-diameter holes machined to depths of 0.5 in. into its outer (curved) surface. The areal density of the assembly ranged from 62.38 g/cm² (along the centerline) to 99.9 g/cm² (along the limb of the polyethylene insert).

Reconstructions of this object using both filtered backprojection and our CCG algorithm with non-negativity constraints are shown in **Fig. 8**. The superiority of the CCG reconstruction is evident. Although we used the squared error criterion with CCG, in fact all the raw data acquired during the OUAL experiments initially bore random sharp spikes rising several hundred to several thousand counts above the local average.

These spikes were due primarily to cosmic ray strikes in the CCD detector used to collect the data. Currently these spikes are removed by preprocessing prior to applying the tomography reconstruction algorithms. However, such data is a perfect candidate

Figure 7. Neutron imaging object.



for the robust techniques we developed this year, and we plan to apply them to neutron data in the near future. For a detailed description of the neutron imaging experiments, see Reference 27.

The Advanced Hydrotest Facility

Last year we adapted our CCG algorithm to cone beam tomography problems. This year we used the resulting algorithm, CCG_Cone, extensively, to study limited-view reconstruction for LLNL's AHF. This code has been an indispensable tool in these studies. It has been used to study the efficacy of reconstruction as a function of number of views, as well as how the orientation of views affects the quality of the reconstruction.

These studies have allowed the AHF design group to make recommendations on how many views are needed, as well as where they should be placed. In addition, CCG_Cone has been used to study the effect of constraints on the quality of reconstructions. It has been found that a judicious use of constraints will help the AHF to achieve its objectives. More programmatic work will continue in this area.

Pulsed Photothermal Radiography

In collaboration with researchers at other institutions, we continued to apply our optimization algorithms to the problem of pulsed photothermal radiography (PPTR). This is essentially a tomographic

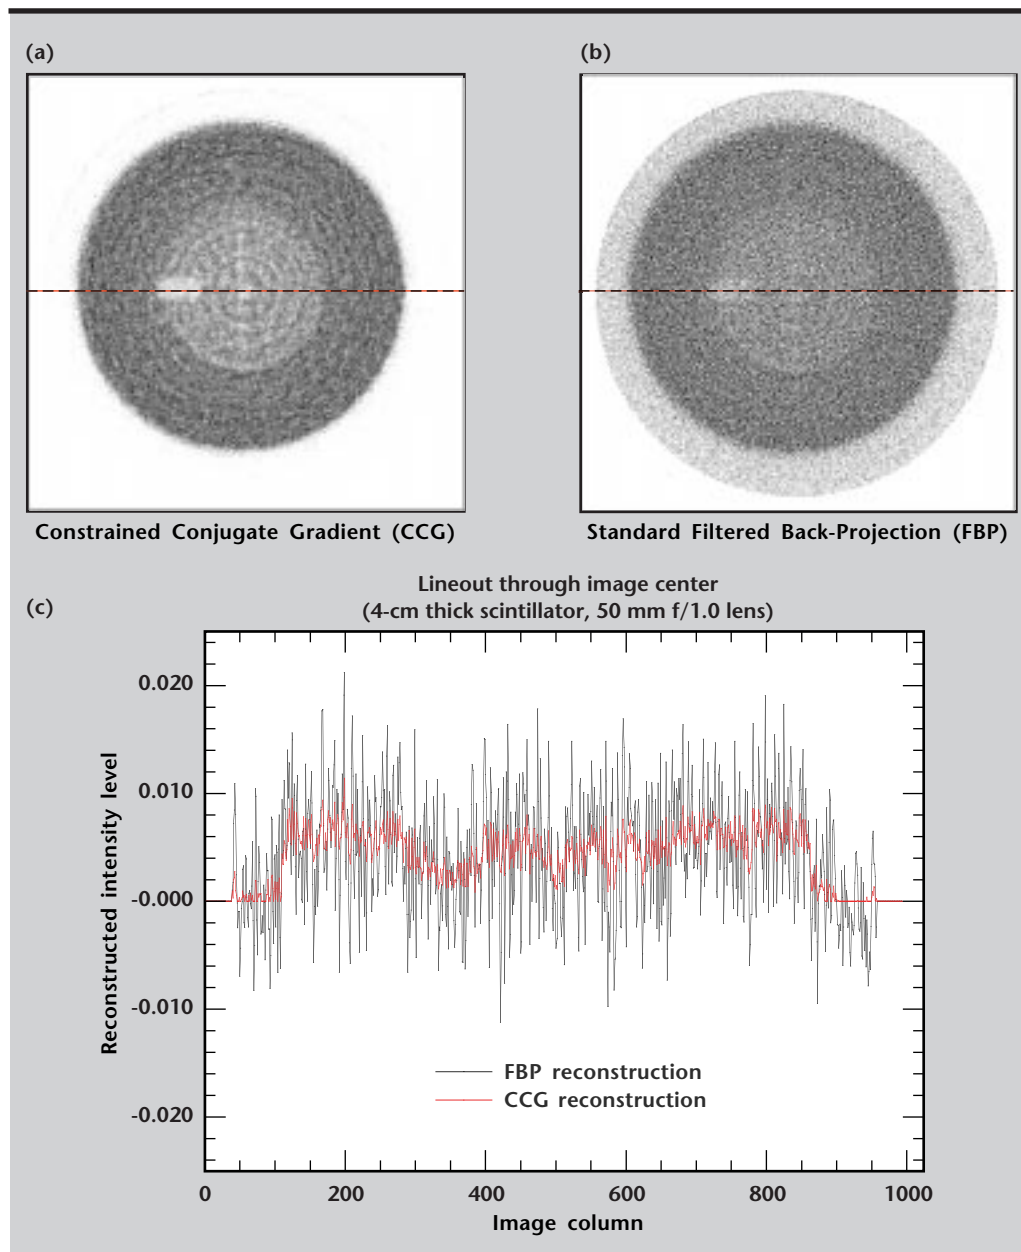


Figure 8.
Reconstruction of
object from neutron
imaging data, using
CCG (a) and filtered
back-projection (b).

method that inverts data from the time evolution of the heat equation, rather than from x-ray projection data, to see inside an opaque object. This year we obtained the first high-quality 3-D images of portwine stain blood vessels,^{28,29} and we demonstrated the feasibility of parallelizing our CCG algorithm.³⁰

Future Work

In future work we plan to continue refining our algorithms and applying them to practical problems. Immediate plans include parallelizing the algorithms, particularly the implementations of **Eq. 1**; further refinements of the forward projection model, including implementing the AHF blurring model; and further investigations into applying the generalized linear model formalism to tomography.

References

1. Poor, H. V. (1994), *An Introduction to Signal Detection and Estimation*, 2nd ed., Springer-Verlag.
2. Gill, P. E., W. Murray, and M. H. Wright (1981), *Practical Optimization*, Academic Press, New York, New York.
3. Fletcher, R. (1987), *Practical Methods of Optimization*, 2nd ed., John C. Wiley and Sons, New York, New York.
4. Luenberger, D. G. (1984), *Linear and Nonlinear Programming*, 2nd ed., Addison-Wesley, Reading, Massachusetts, pp. 214–220.
5. Stark, H., ed. (1987), *Image Recovery: Theory and Application*, Academic Press, New York, New York.
6. Azevedo, S. G. (1991), Model-Based Computed Tomography for Nondestructive Evaluation, PhD Thesis, Lawrence Livermore National Laboratory, Livermore, California (UCRL-LR-106884), March.
7. Goodman, D. M., E. M. Johansson, and T. W. Lawrence (1993), "On Applying the Conjugate Gradient Algorithm to Image Processing Problems," Chapter 11 in *Multivariate Analysis: Future Directions*, C. R. Rao, ed., Elsevier Science Publishers.
8. Goodman, D. M., T. W. Lawrence, E. M. Johansson, and J. P. Fitch (1993), "Bispectral Speckle Interferometry to Reconstruct Extended Objects from Turbulence-Degraded Telescope Images," Chapter 13 in *Handbook of Statistics, Vol. 10: Signal Processing and its Applications*, N. K. Bose and C. R. Rao, eds., North Holland, Amsterdam.
9. Kolman, J., W. S. Haddad, D. M. Goodman, and K. A. Nugent (1994), "Application of a Constrained Optimization Algorithm to Limited View Tomography," *Proc. SPIE Conf.*, San Diego, California, July.
10. Haddad, W. S., J. E. Trebes, D. M. Goodman, H. R. Lee, I. McNulty, E. H. Anderson, and A. O. Zalensky (1995), "Ultra High Resolution Soft X-Ray Tomography," *Proc. SPIE Conf.*, San Diego, California, July.
11. Somoza, J. R., H. Szöke, D. M. Goodman, P. Beran, D. Truckses, S.-H. Kim, and A. Szöke (1995), "Holographic Methods in X-ray Crystallography IV. A Fast Algorithm and its Application to Macromolecular Crystallography," *Acta Crystallographica*, Vol. **A51**.
12. Milner, T. E., D. M. Goodman, B. S. Tanenbaum, and J. S. Nelson (1995), "Depth Profiling of Laser-Heated Chromophores in Biological Tissues by Pulsed Photothermal Radiometry," *Journal of the Optical Society of America-A*, Vol. **12** No. 7, July.
13. Milner, T. E., D. M. Goodman, B. S. Tanenbaum, B. Anvari, L. O. Svasand, and J. S. Nelson (1996), "Imaging of Laser Heated Subsurface Chromophores in Biological Materials: Determination of Lateral Physical Dimensions," *Physics in Medicine and Biology*, Vol. **41**.
14. Milner, T. E., D. J. Smithies, D. M. Goodman, A. Lau, and J. S. Nelson (1996), "Depth Determination of Chromophores in Human Skin by Pulsed Photothermal Radiometry," *Applied Optics*, June.
15. van Gemert, M. J. C., J. S. Nelson, T. E. Milner, D. J. Smithies, W. Verkruijsse, J. F. de Boer, G. W. Lucassen, D. M. Goodman, B. S. Tanenbaum, L. T. Norvang, and L. O. Svasand (1997), "Non-Invasive Determination of Port Wine Stain Anatomy and Physiology for Optimal Laser Treatment Strategies," *Physics in Medicine and Biology*, Vol. **42**.
16. Byrd, R. H., J. Nocedal, and R. B. Schnabel (1994), "Representations of Quasi-Newton Matrices and Their Use in Limited Memory Methods," *Math. Prog.*, Vol. **63**, No. 2, January.
17. "Advanced Hydrodynamic Radiography Technology Development Plan" (1996), Predecisional Draft, US Department of Energy, Defense Programs, Office of Research and Inertial Fusion, February 16.
18. Huber, P. J. (1981), *Robust Statistics*, John C. Wiley and Sons, New York, New York.
19. Myers, R. H. (1990), *Classical and Modern Regression with Applications*, PWS-Kent, Boston, Massachusetts, Chapter 7.
20. Martz, H. E., G. P. Roberson, D. C. Camp, D. J. Decman, J. A. Jackson, and G. K. Becker (1998), "Active and Passive Computed Tomography Mixed Waste Area Final Report," Lawrence Livermore National Laboratory, Livermore, California, November.

21. Brown, J. K., K. Kalki, J. A. Heanue, and B. H. Hasegawa (1995), "Quantitative SPECT Reconstruction Using Multiray Projection Integrators," *Conference Record 1995 IEEE Nuclear Science Symposium and Medical Imaging Conference*, Vol. 2, pp. 1272–1276.
22. Keto, E., S. G. Azevedo, G. P. Roberson, D. J. Decman, H. E. Martz, and E. M. Johansson (1995), "Spatial Resolution Versus Signal to Noise in Quantitative Tomography," *Proceedings of the 4th Nondestructive Assay and Nondestructive Examination Waste Characterization Conference*, Salt Lake City, Utah, October 24–26, pp. 405–420.
23. Goodman, D. M. (1997), "Maximum Likelihood Estimation with Poisson (Counting) Statistics for Waste Drum Inspection," Lawrence Livermore National Laboratory, Livermore, California (UCRL-ID-127361), May.
24. Jackson, J. A., D. M. Goodman, G. P. Roberson, and H. E. Martz (1998), "An Active and Passive Computed Tomography Algorithm with a Constrained Conjugate Gradient Solution," *6th Nondestructive Assay Waste Characterization Conference*, Salt Lake City, Utah, November 17–19.
25. McCullagh, P., and J. A. Nelder (1989), *Generalized Linear Models*, 2nd ed., Chapman and Hall, London.
26. Dobson, A. J. (1990), *An Introduction to Generalized Linear Models*, Chapman and Hall, London.
27. Hall, J., F. Dietrich, C. Logan, and G. Schmid (1998), "Development of High-Energy Neutron Imaging in Support of Enhanced Surveillance Program Applications," Lawrence Livermore National Laboratory, Livermore, California, in preparation.
28. Shoari, S., N. Bagherzadeh, D. Goodman, T. E. Milner, D. J. Smithies, and J. S. Nelson (1998), "A Parallel Algorithm for Pulsed Laser Infrared Tomography," *Pattern Recognition Letters*, (19)5-6, pp. 521–526, April.
29. Milner, T. E., S. A. Telenkov, B. S. Tanenbaum, J. S. Nelson, and D. M. Goodman (1998), "Non-invasive Evaluation of Biological Materials Using Pulsed Photothermal Tomography," *Proc. Winter Annual Meeting of ASME*, Anaheim, California, November.
30. Telenkov, S. A., D. M. Goodman, B. S. Tanenbaum, J. S. Nelson, and T. E. Milner (1998), "Infrared Imaging of laser-heated port wine stains," *Annual Meeting of the Optical Society of America*, Baltimore, Maryland, October. 